



(12) **United States Patent**  
**Raisch et al.**

(10) **Patent No.:** **US 11,036,661 B2**  
(45) **Date of Patent:** **Jun. 15, 2021**

(54) **DIRECTED INTERRUPT VIRTUALIZATION**

(56) **References Cited**

(71) Applicant: **INTERNATIONAL BUSINESS MACHINES CORPORATION**, Armonk, NY (US)  
(72) Inventors: **Christoph Raisch**, Gerlingen (DE); **Marco Kraemer**, Sindelfingen (DE); **Donald William Schmidt**, Stone Ridge, NY (US); **Bernd Nerz**, Boeblingen (DE); **Frank Siegfried Lehnert**, Weil im Schoenbuch (DE); **Peter Dana Driever**, Poughkeepsie, NY (US)

U.S. PATENT DOCUMENTS

5,577,255 A 11/1996 Roemer  
7,447,820 B2 11/2008 Raj  
(Continued)

FOREIGN PATENT DOCUMENTS

TW 1537877 B 11/2016  
TW 1570563 B 11/2017

OTHER PUBLICATIONS

IBM, "z/Architecture—Principles of Operation," IBM Publication No. SA22-7832-12, Thirteenth Edition, Sep. 2019, pp. 1-2000.  
(Continued)

(73) Assignee: **INTERNATIONAL BUSINESS MACHINES CORPORATION**, Armonk, NY (US)

(\* ) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

*Primary Examiner* — Henry Tsai  
*Assistant Examiner* — Christopher A Daley  
(74) *Attorney, Agent, or Firm* — Steven Chiu, Esq.; Blanche E. Schiller, Esq.; Heslin Rothenberg Farley & Mesiti P.C.

(21) Appl. No.: **16/789,504**

(22) Filed: **Feb. 13, 2020**

(65) **Prior Publication Data**  
US 2020/0264992 A1 Aug. 20, 2020

(57) **ABSTRACT**

An interrupt signal is provided to a guest operating system executed using one or more processors of a plurality of processors. One or more bus connected modules are operationally connected with the plurality of processors via a bus and a bus attachment device. The bus attachment device receives an interrupt signal from one of the bus connected modules with an interrupt target ID identifying one of the processors assigned for use by the guest operating system as a target processor for handling the interrupt signal. The bus attachment device translates the received interrupt target ID to a logical processor ID of the target processor using a mapping table comprised by the bus attachment device and forwards the interrupt signal to the target processor for handling. The logical processor ID of the target processor is used to address the target processor directly.


(30) **Foreign Application Priority Data**  
Feb. 14, 2019 (EP) ..... 19157093

(51) **Int. Cl.**  
**G06F 13/24** (2006.01)  
**G06F 9/455** (2018.01)  
(Continued)

(52) **U.S. Cl.**  
CPC ..... **G06F 13/24** (2013.01); **G06F 9/45558** (2013.01); **G06F 13/1668** (2013.01); **G06F 13/28** (2013.01)

(58) **Field of Classification Search**  
CPC ..... G06F 13/24; G06F 13/26; G06F 13/1668; G06F 9/45558  
See application file for complete search history.

**20 Claims, 23 Drawing Sheets**



- (51) **Int. Cl.**  
**G06F 13/16** (2006.01)  
**G06F 13/28** (2006.01)

OTHER PUBLICATIONS

(56) **References Cited**

U.S. PATENT DOCUMENTS

8,453,143	B2	5/2013	Mahalingam	
8,601,194	B2	12/2013	Tsirkin	
9,032,127	B2	5/2015	Fischer	
9,043,521	B2	5/2015	Tiruvallur	
9,378,162	B2 *	6/2016	Jebson	G06F 9/4812
9,465,760	B2 *	10/2016	Egi	G06F 13/32
9,507,740	B2 *	11/2016	Feehrer	G06F 13/4068
9,697,029	B2	7/2017	Tsirkin	
9,772,868	B2	9/2017	Tu	
9,830,286	B2 *	11/2017	Tsirkin	G06F 13/24
9,910,699	B2 *	3/2018	Sankaran	G06F 13/24
9,910,700	B2	3/2018	Prosch	
9,952,987	B2	4/2018	Guddeti et al.	
10,078,603	B2 *	9/2018	Tsirkin	G06F 13/24
10,713,095	B2 *	7/2020	Yoo	G06F 9/5044
2003/0200250	A1	10/2003	Klick	
2003/0204655	A1	10/2003	Schmisser	
2005/0080982	A1	4/2005	Vasilevsky	
2005/0125580	A1	6/2005	Madukkarumukumana	
2006/0233203	A1	10/2006	Iwamura	
2011/0047309	A1 *	2/2011	Brinkmann	G06F 13/24 710/261
2011/0197003	A1	8/2011	Serebrin et al.	
2011/0273309	A1	11/2011	Zhang	
2011/0320664	A1	12/2011	Belmar et al.	
2011/0320840	A1 *	12/2011	Nayar	G06F 9/5094 713/323
2012/0099521	A1	4/2012	Ryu	
2013/0080673	A1	3/2013	Feehrer	
2013/0104127	A1	4/2013	Domsch	
2013/0159579	A1	6/2013	Neiger et al.	
2014/0047150	A1	2/2014	Marietta	
2014/0223060	A1	8/2014	Tsirkin et al.	
2015/0312116	A1	10/2015	Taheri	
2016/0117190	A1	4/2016	Sankaran et al.	
2018/0129619	A1	5/2018	Neiger et al.	
2018/0356964	A1	12/2018	Morris	
2019/0163519	A1 *	5/2019	Freeman	G06F 9/4818
2020/0264910	A1	8/2020	Kraemer	
2020/0264911	A1	8/2020	Raisch	
2020/0264912	A1	8/2020	Raisch	
2020/0264917	A1	8/2020	Kraemer	
2020/0264993	A1	8/2020	Nerz	
2020/0264994	A1	8/2020	Raisch	
2020/0264995	A1	8/2020	Raisch	

Mell, Peter and Tim Grance, "The NIST Definition of Cloud Computing," National Institute of Standards and Technology, Information Technology Laboratory, Special Publication 800-145, Sep. 2011, pp. 1-7.

Intel 64 and IA-32 Architectures Software Developer's Manual—vol. 3A; System Programming Guide, Part 1, Sep. 2016, Chapter 10 (10-1-10-50) (+ cover pages).

Williams, Chris, "Intel's Broadwell Xeon E5-2600 v4 chips: So What's in it for you, smartie-pants coders," Mar. 2016, pp. 1-9.

Raj, Ashok, "Recent Enhancements in Intel Virtualization Technology for Directed I/O (Intel VT-D)," Sep. 2018, pp. 1-7.

Kraemer, Marco et al., "Directed Interrupt Virtualization With Interrupt Table," U.S. Appl. No. 16/789,515, filed Feb. 13, 2020, pp. 1-103.

Kraemer, Marco et al., "Directed Interrupt Virtualization With Running Indicator," U.S. Appl. No. 16/789,519, filed Feb. 13, 2020, pp. 1-105.

Nerz, Bernd et al., "Directed Interrupt for Multilevel Virtualization," U.S. Appl. No. 16/789,532, filed Feb. 13, 2020, pp. 1-109.

Raisch, Christoph et al., "Directed Interrupt for Multilevel Virtualization With Interrupt Table," U.S. Appl. No. 16/789,556, filed Feb. 13, 2020, pp. 1-106.

Raisch, Christoph et al., "Directed Interrupt Virtualization With Blocking Indicator," U.S. Appl. No. 16/789,567, filed Feb. 13, 2020, pp. 1-99.

Raisch, Christoph et al., "Interrupt Signaling for Directed Interrupt Virtualization," U.S. Appl. No. 16/789,581, filed Feb. 13, 2020, pp. 1-101.

Raisch, Christoph et al., "Directed Interrupt Virtualization With Fallback," U.S. Appl. No. 16/789,600, filed Feb. 13, 2020, pp. 1-103.

List of IBM Patents or Patent Applications Treated as Related, May 13, 2020, 2 pages.

Feb. 13, 2020, U.S. Appl. No. 16/789,514, 2020-0264917.

Dec. 30, 2020, U.S. Appl. No. 17/137,755.

Feb. 13, 2020, U.S. Appl. No. 16/789,519, 2020-0264910.

Feb. 13, 2020, U.S. Appl. No. 16/789,532, 2020-0264993.

Feb. 13, 2020, U.S. Appl. No. 16/789,556, 2020-0264994.

Feb. 13, 2020, U.S. Appl. No. 16/789,567, 2020-0264995.

Feb. 13, 2020, U.S. Appl. No. 16/789,581, 2020-0264911.

Nov. 6, 2020, U.S. Appl. No. 17/091,498.

Feb. 13, 2020, U.S. Appl. No. 16/789,600, 2020-0264912.

Raisch, Christoph et al., "Interrupt Signaling for Directed Interrupt Virtualization," U.S. Appl. No. 17/091,498, filed Nov. 6, 2020, pp. 1-98.

Kraemer, Marco et al., "Directed Interrupt Virtualization with Interrupt Table," U.S. Appl. No. 17/137,755, filed Dec. 30, 2020, pp. 1-103.

List of IBM Patents or Patent Applications Treated as Related, Jan. 25, 2021, 2 pages.

\* cited by examiner




FIG. 1




FIG. 2






FIG. 3




FIG. 4




FIG. 5




FIG. 6






FIG. 7

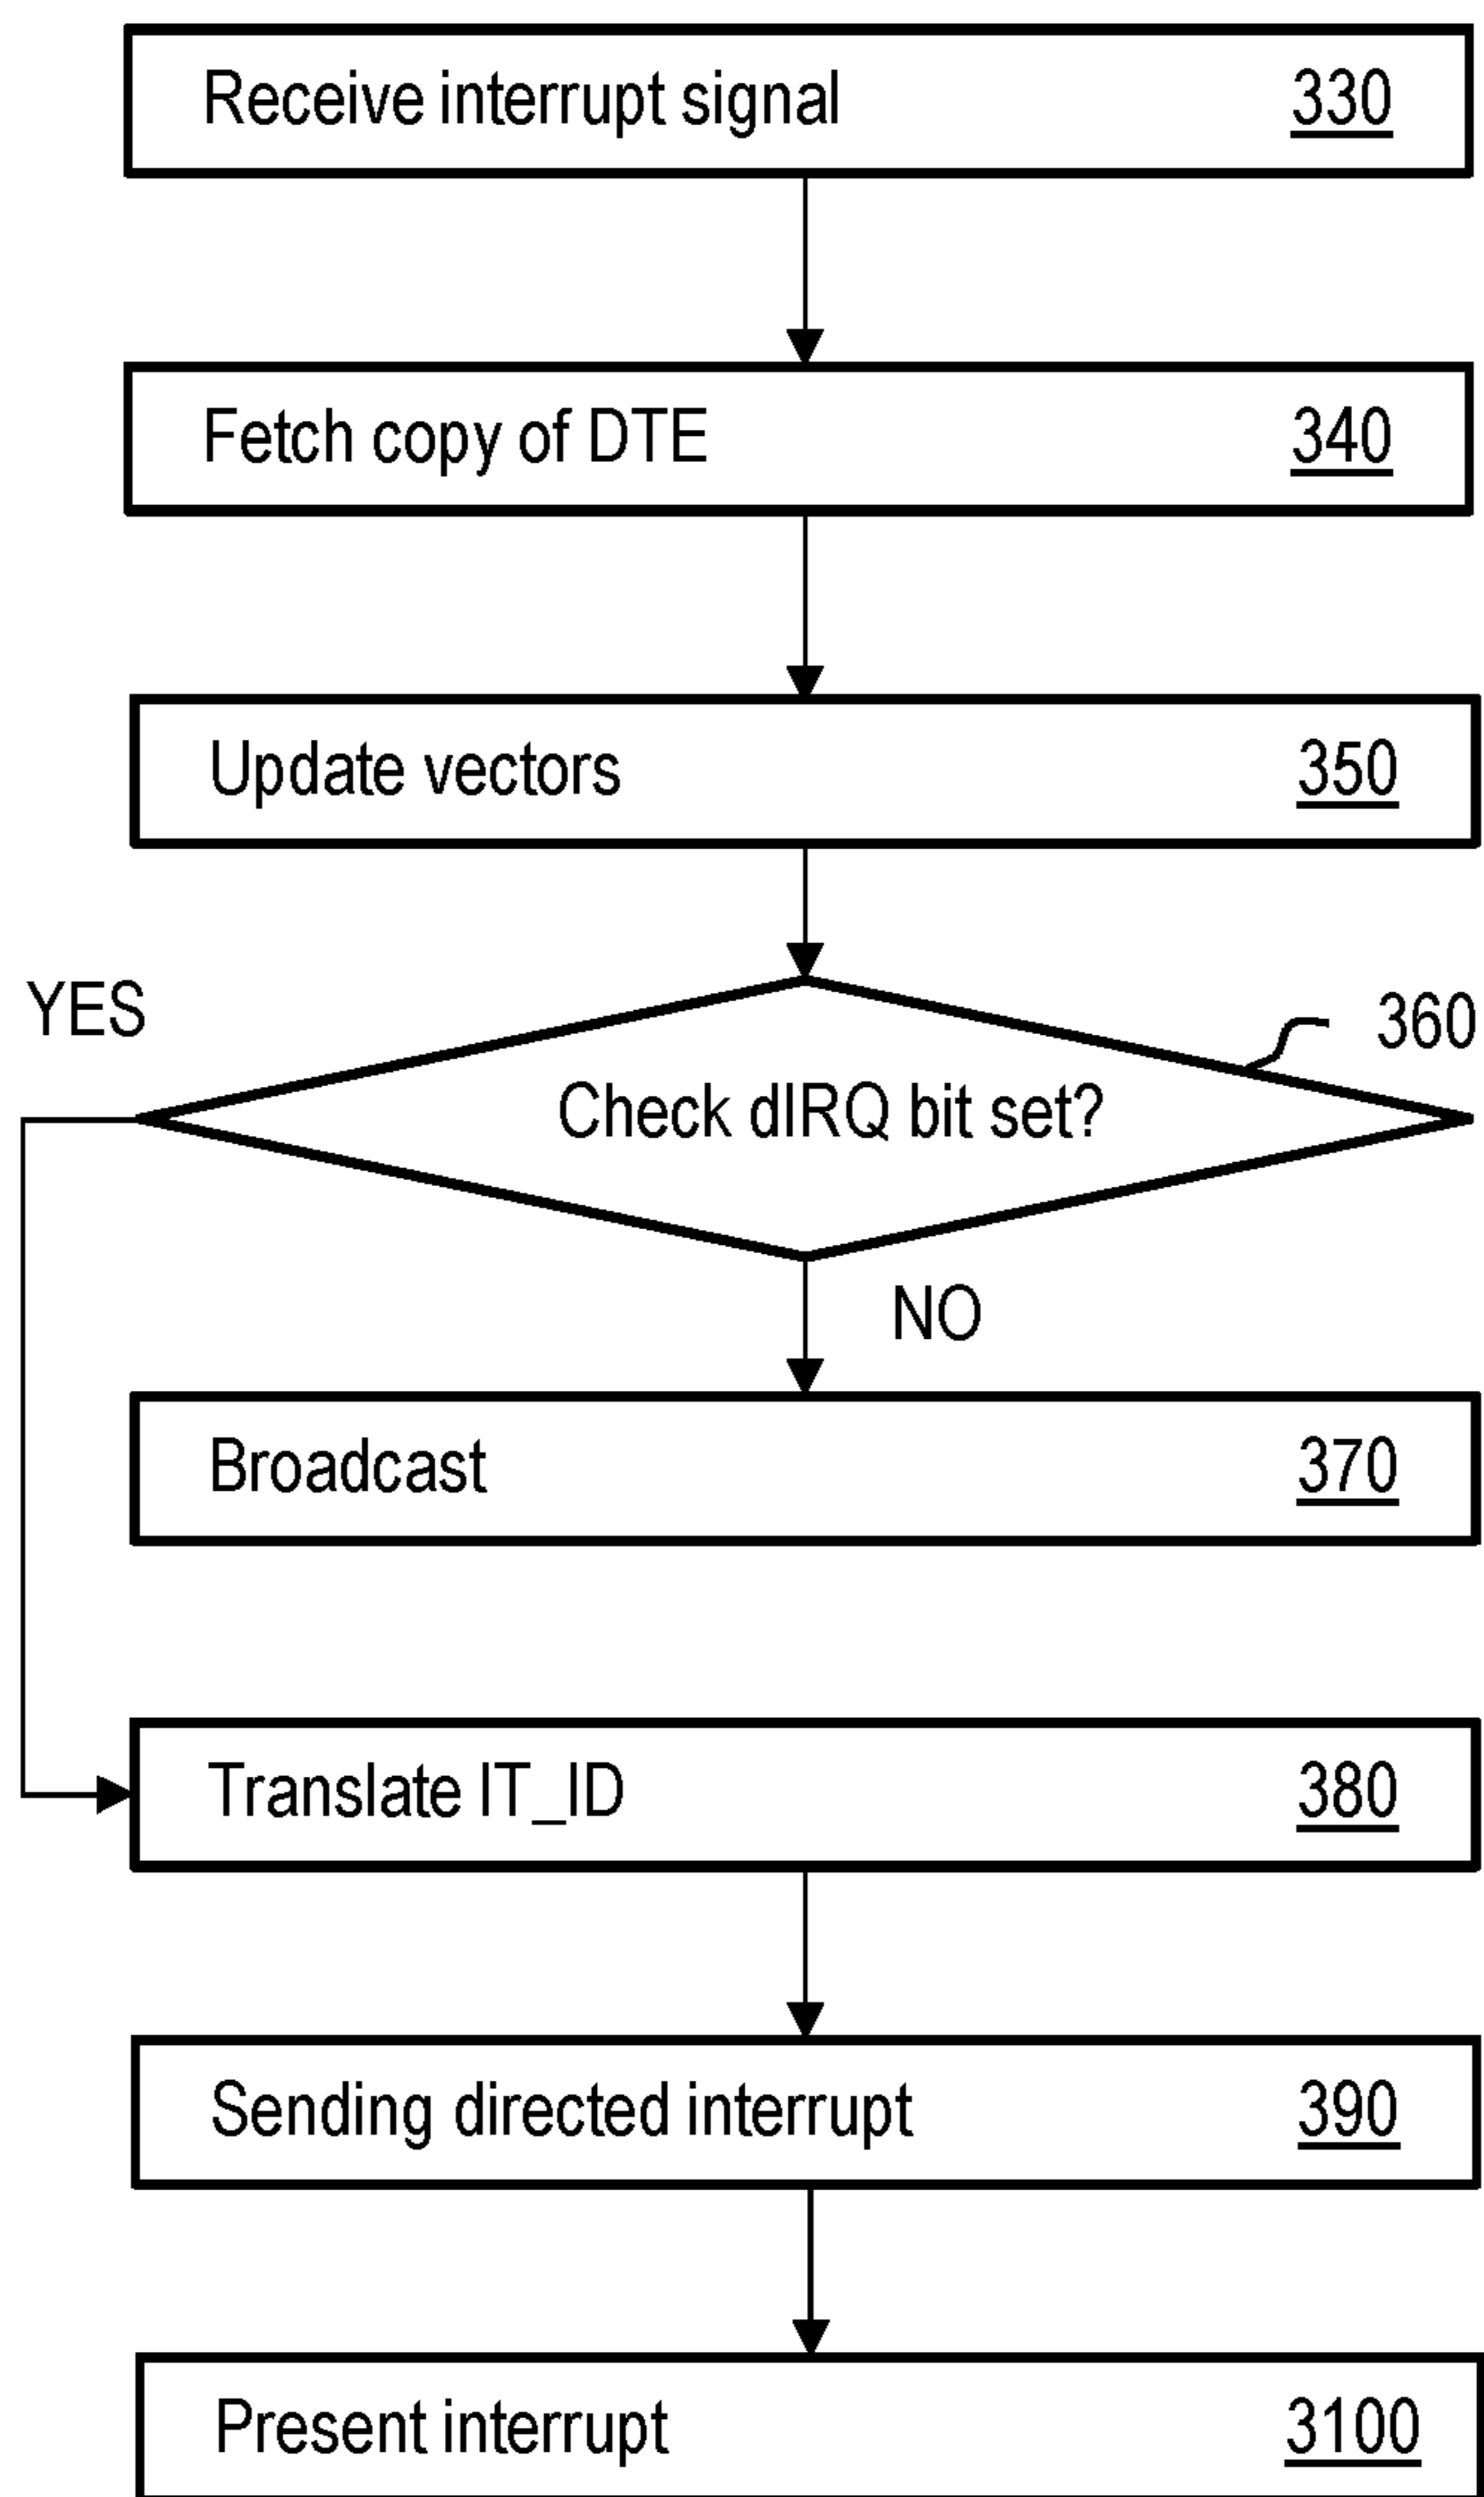


FIG. 8




FIG. 9

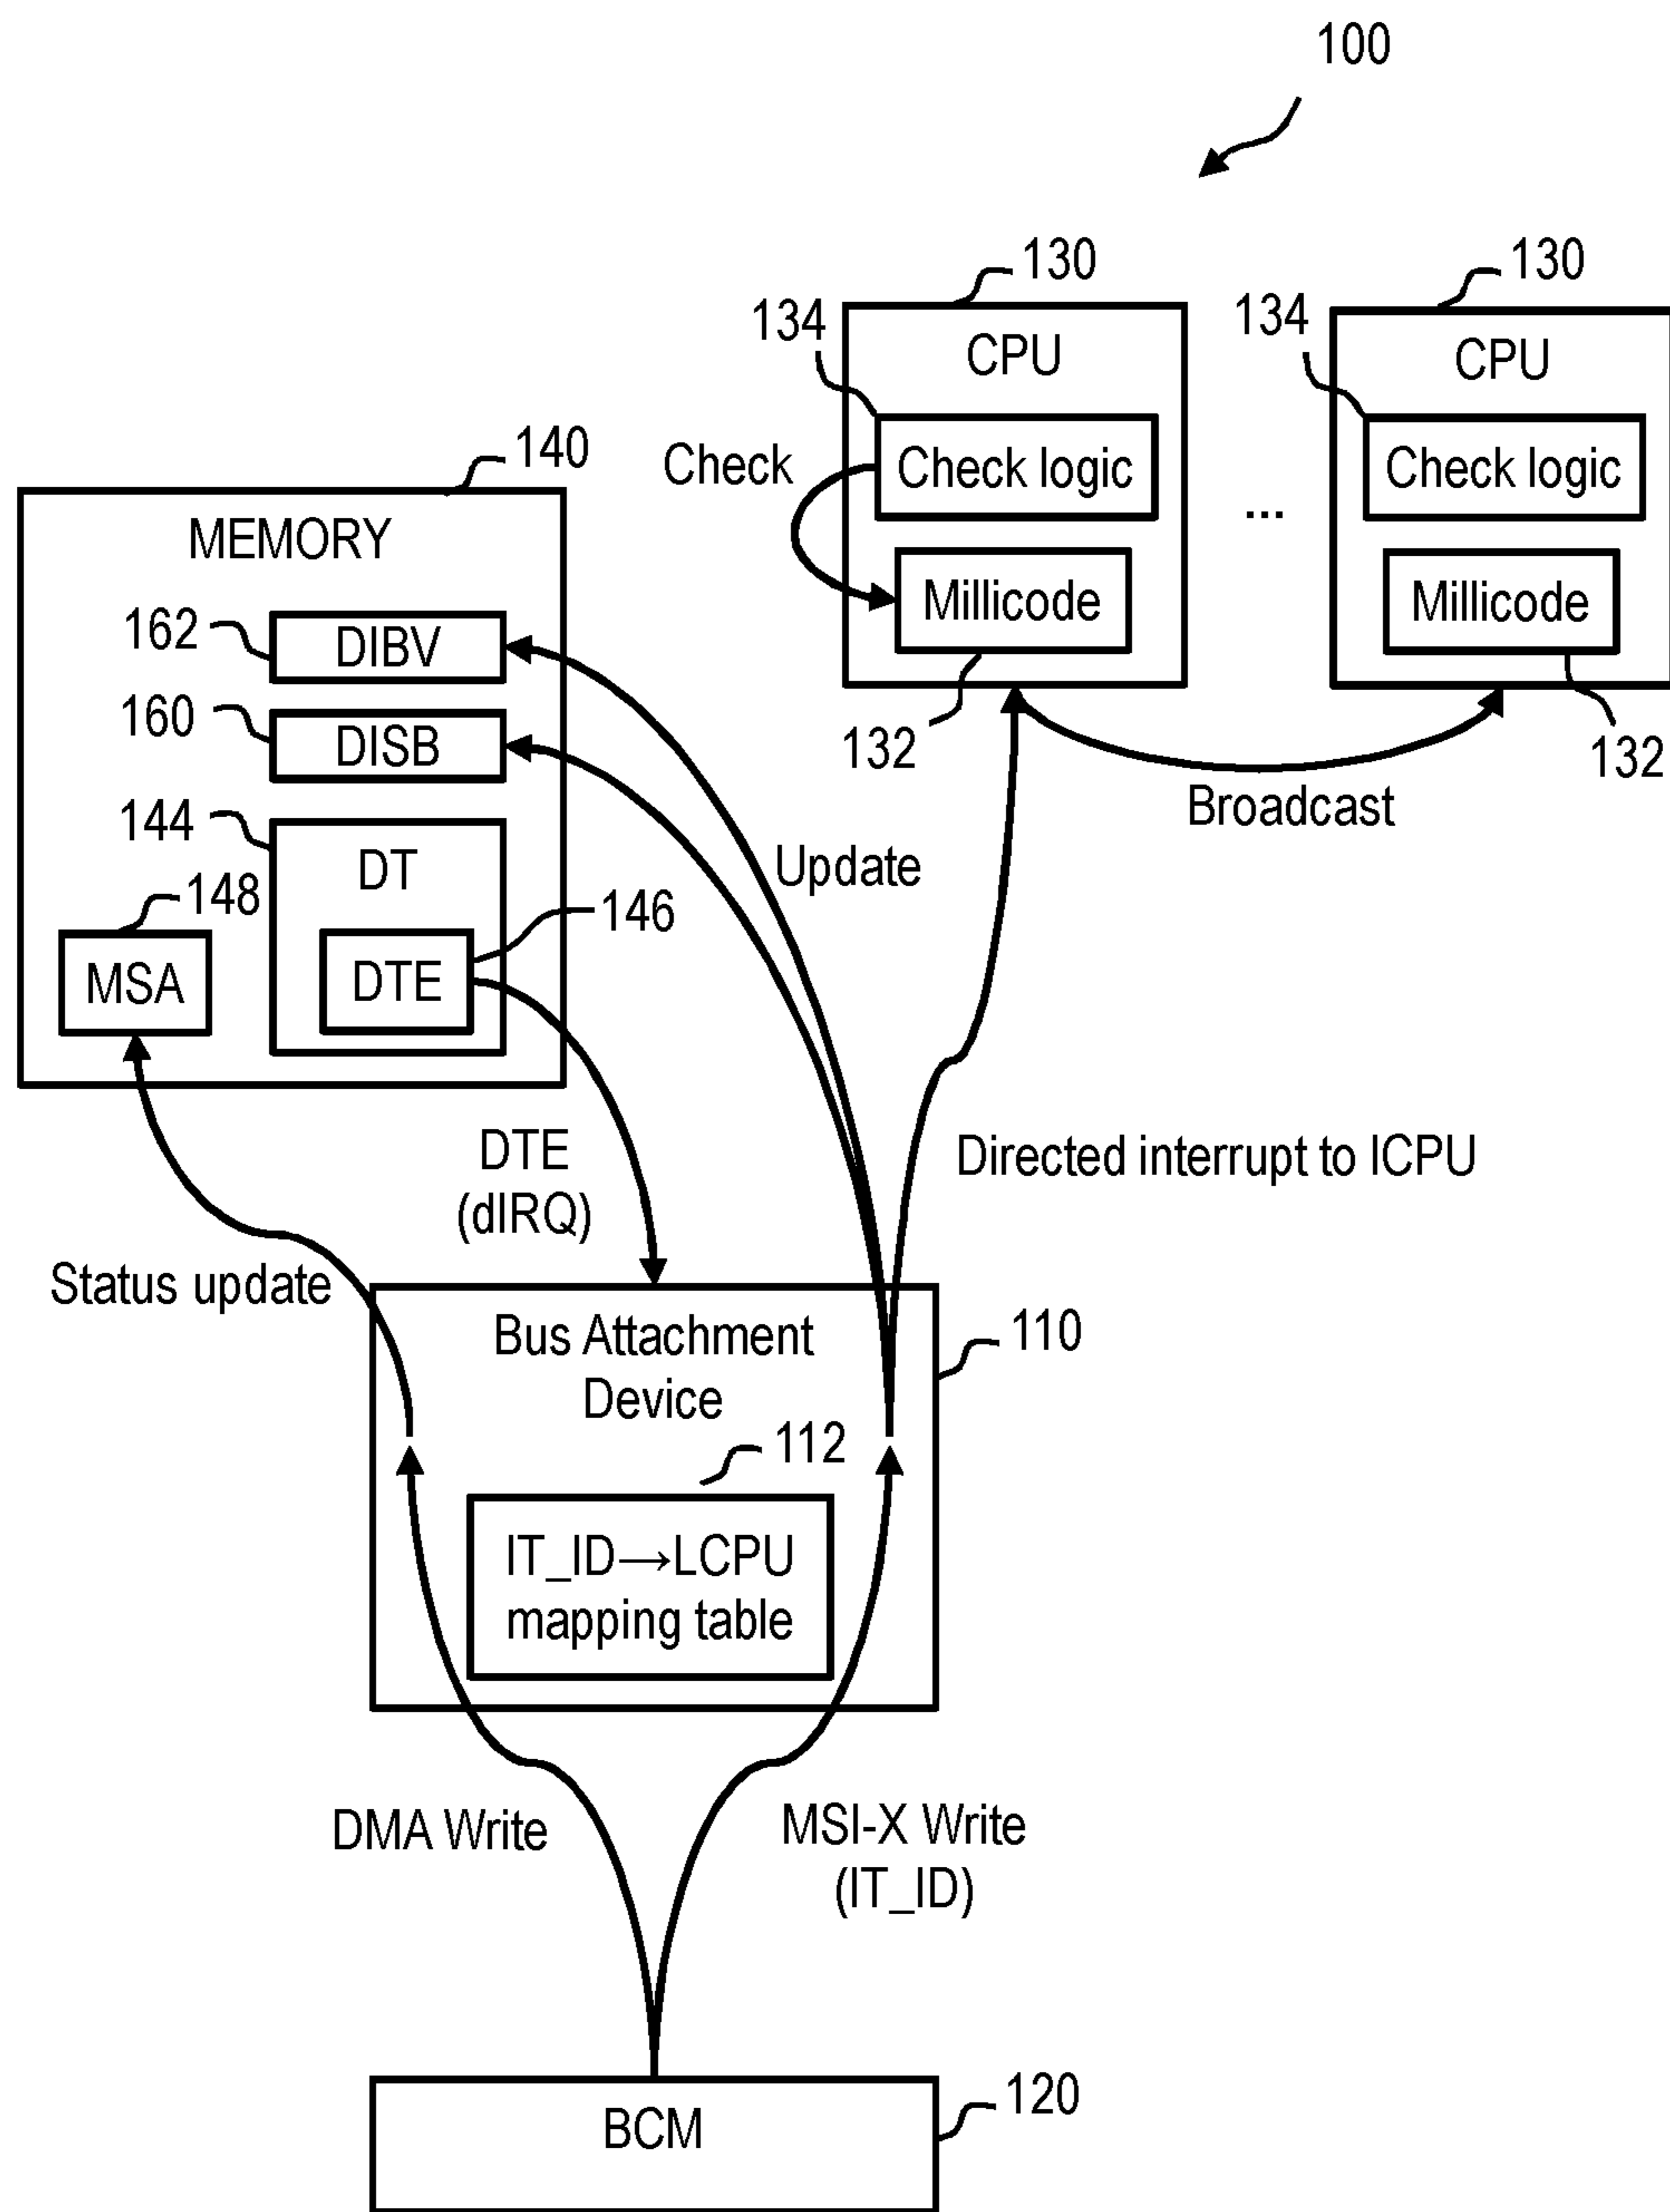


FIG. 10

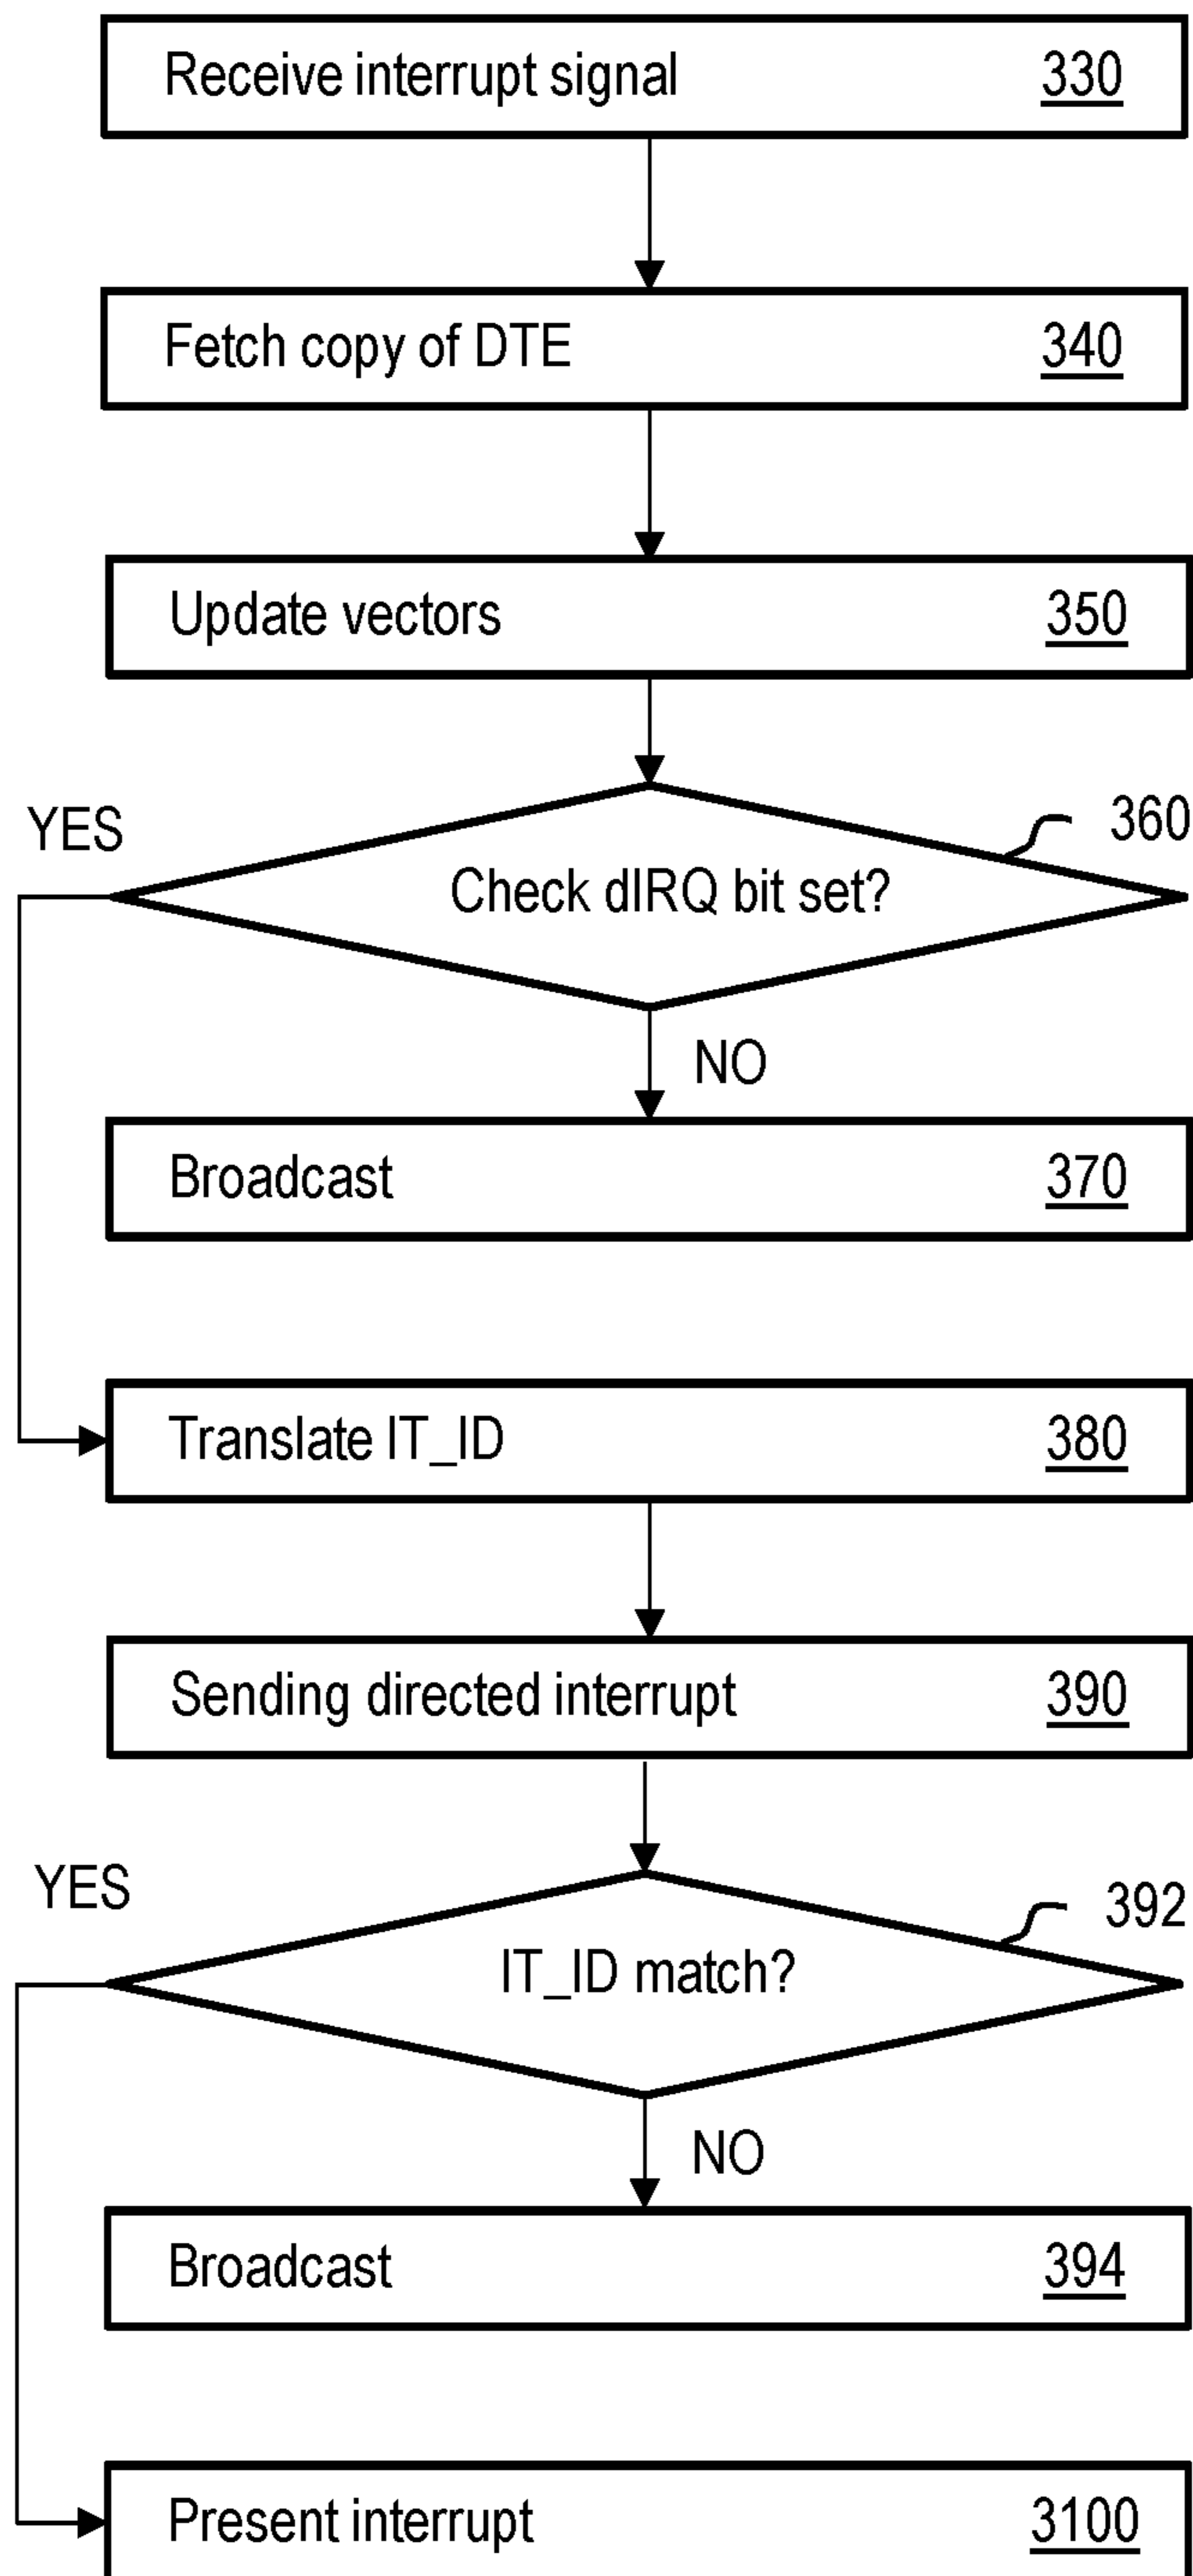


FIG. 11

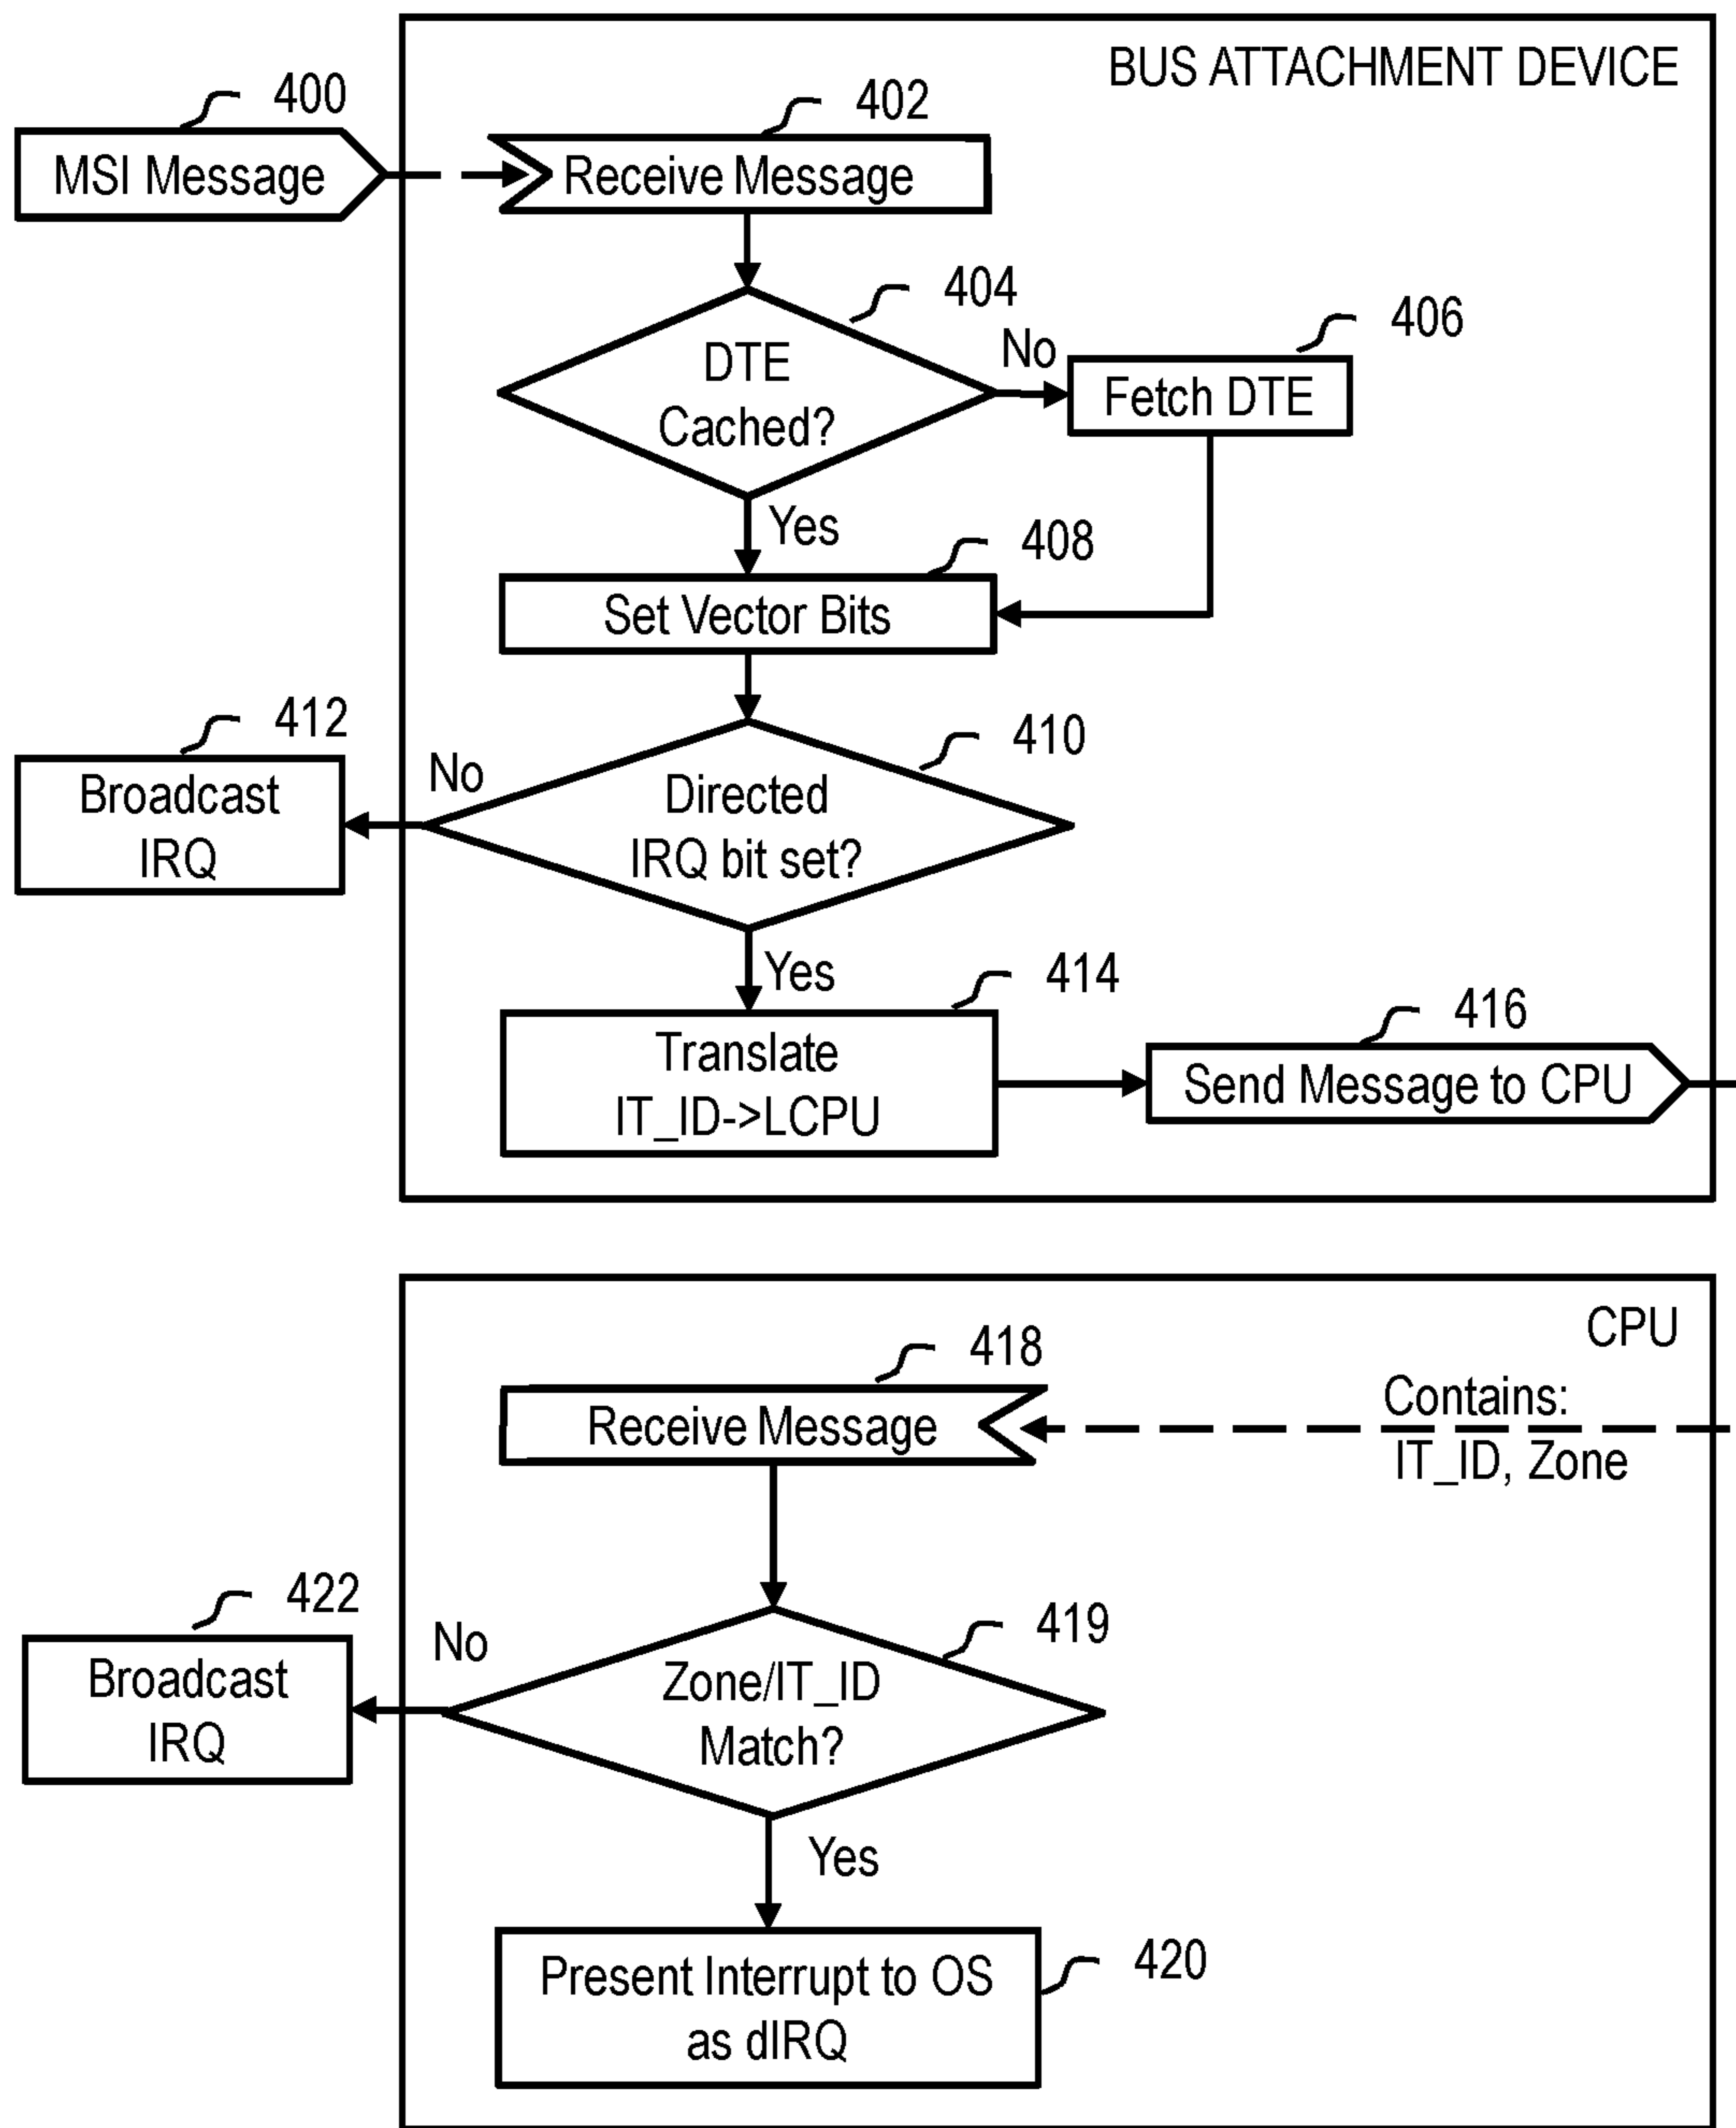


FIG. 12






FIG. 13

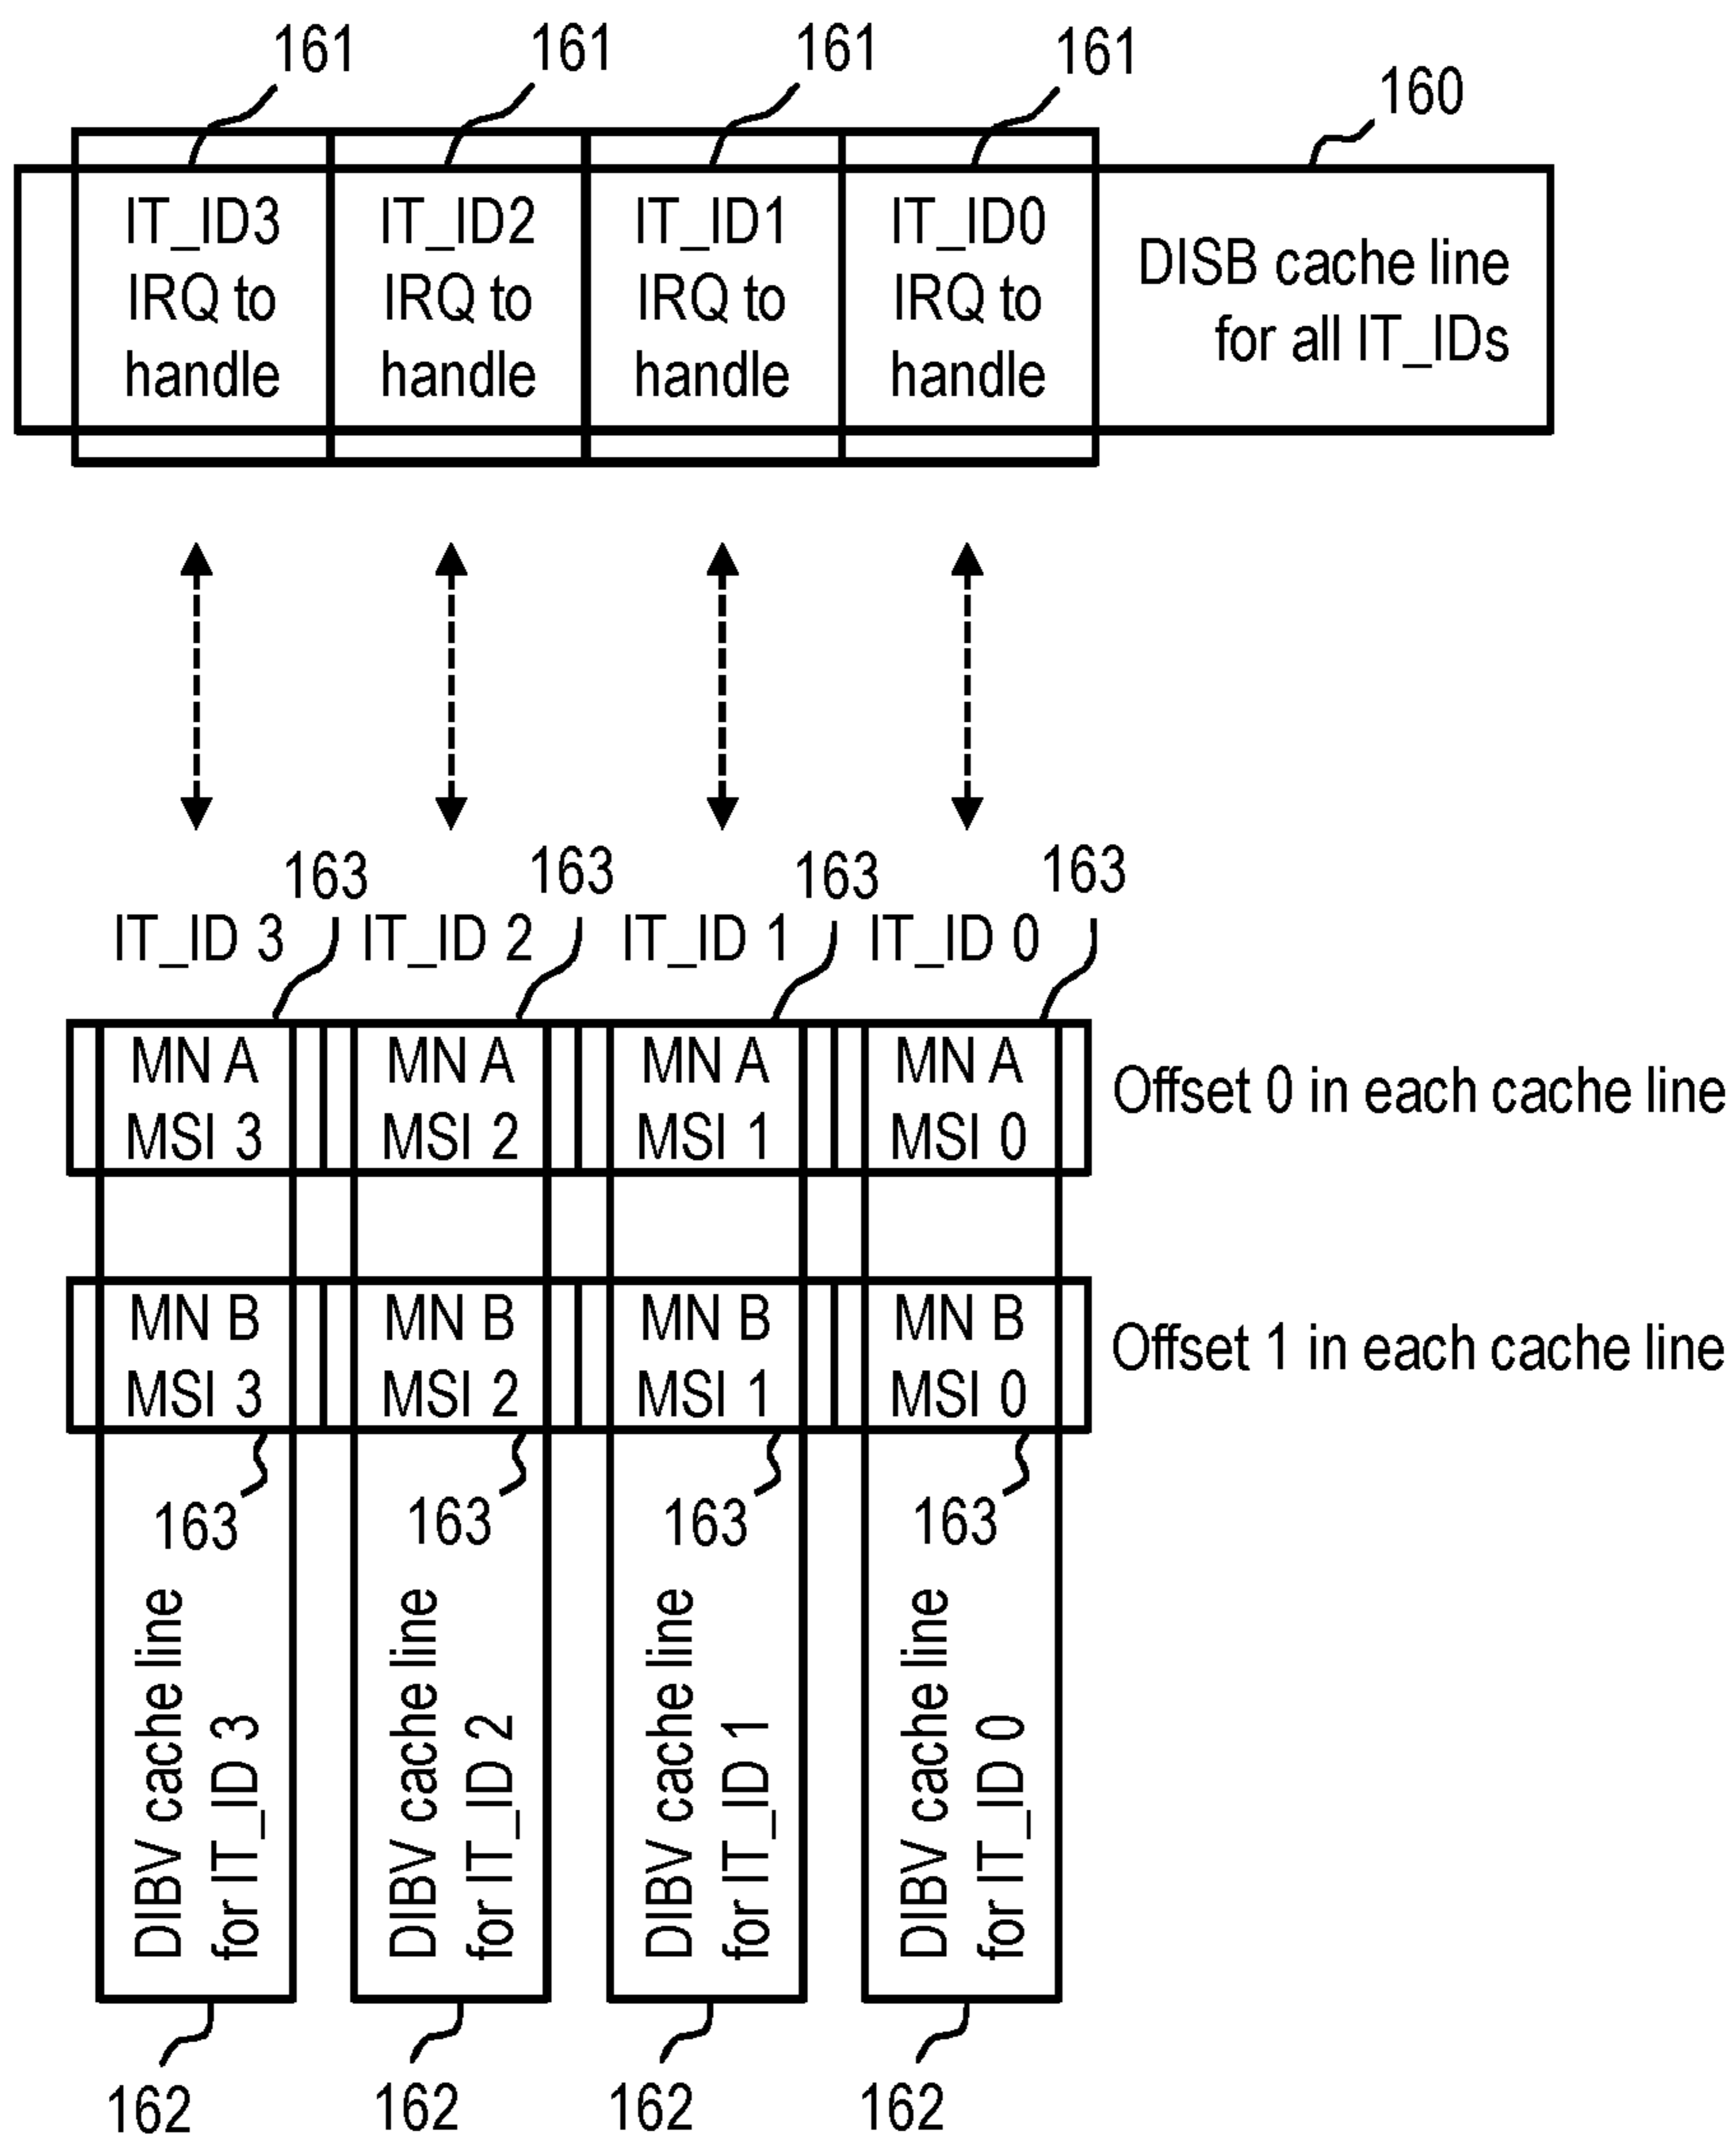


FIG. 14

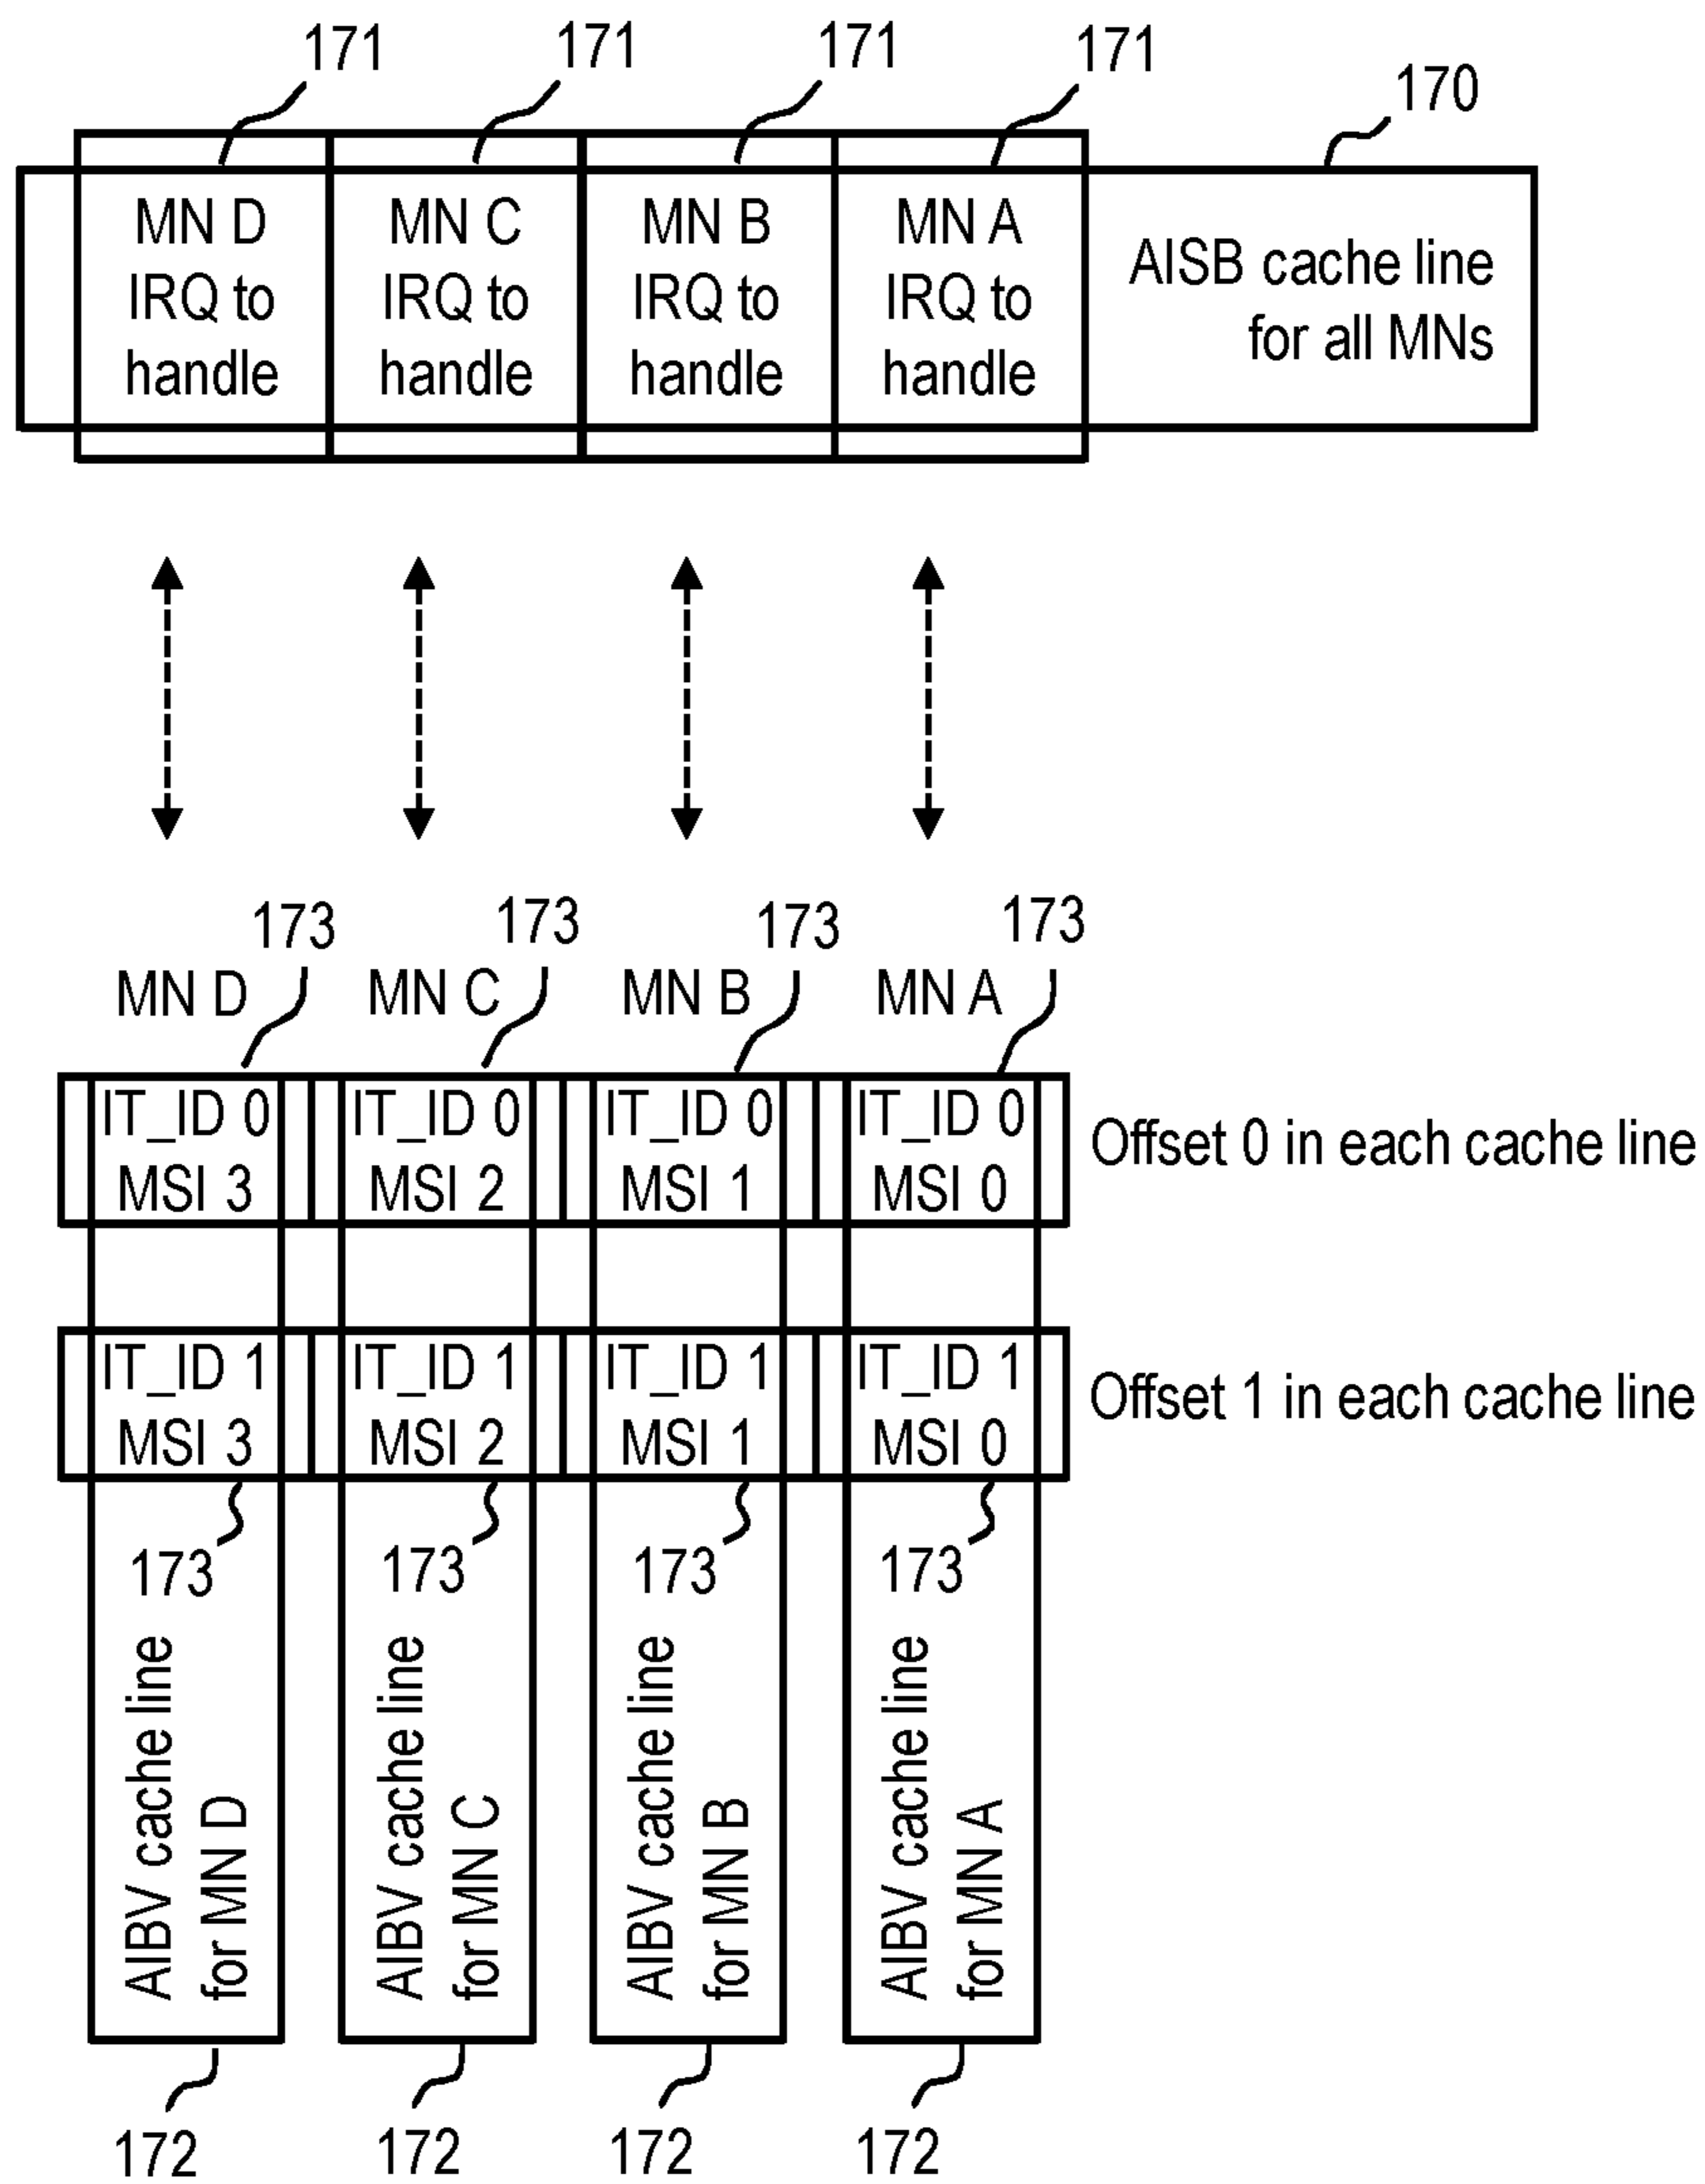


FIG. 15




FIG. 16A




FIG. 16B




FIG. 17A




FIG. 17B




FIG. 18

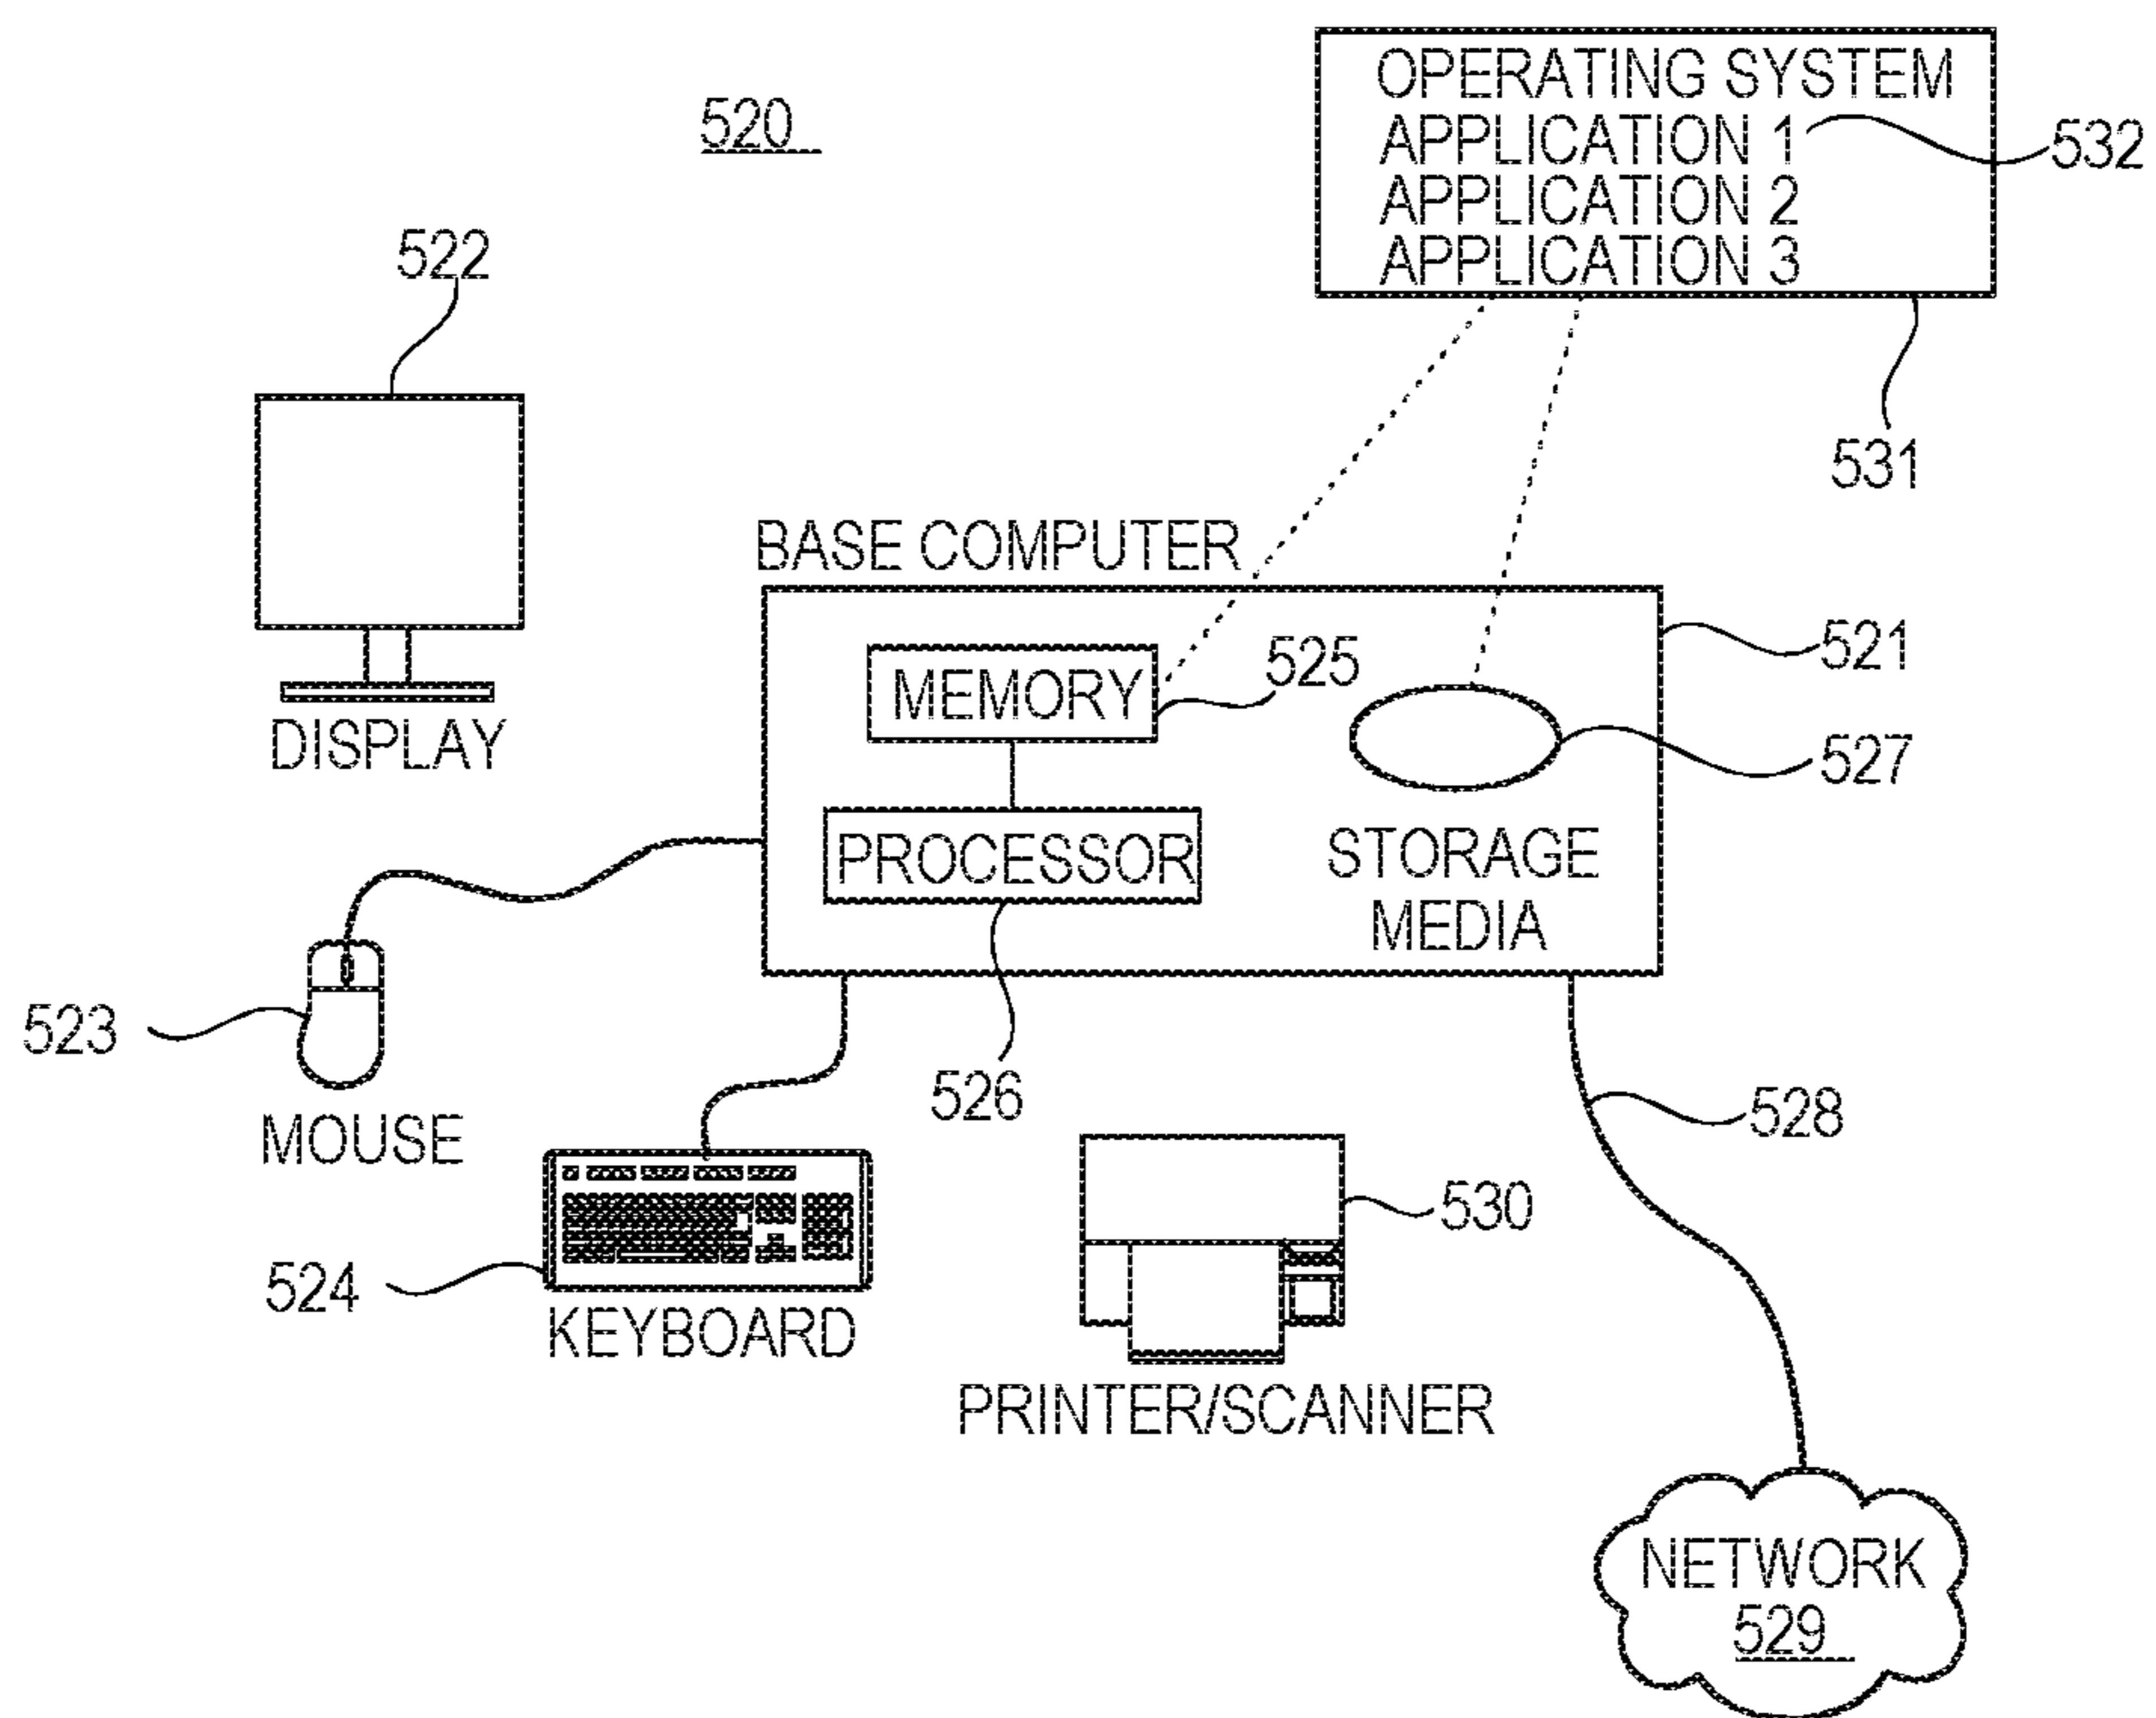


FIG. 19




FIG. 20



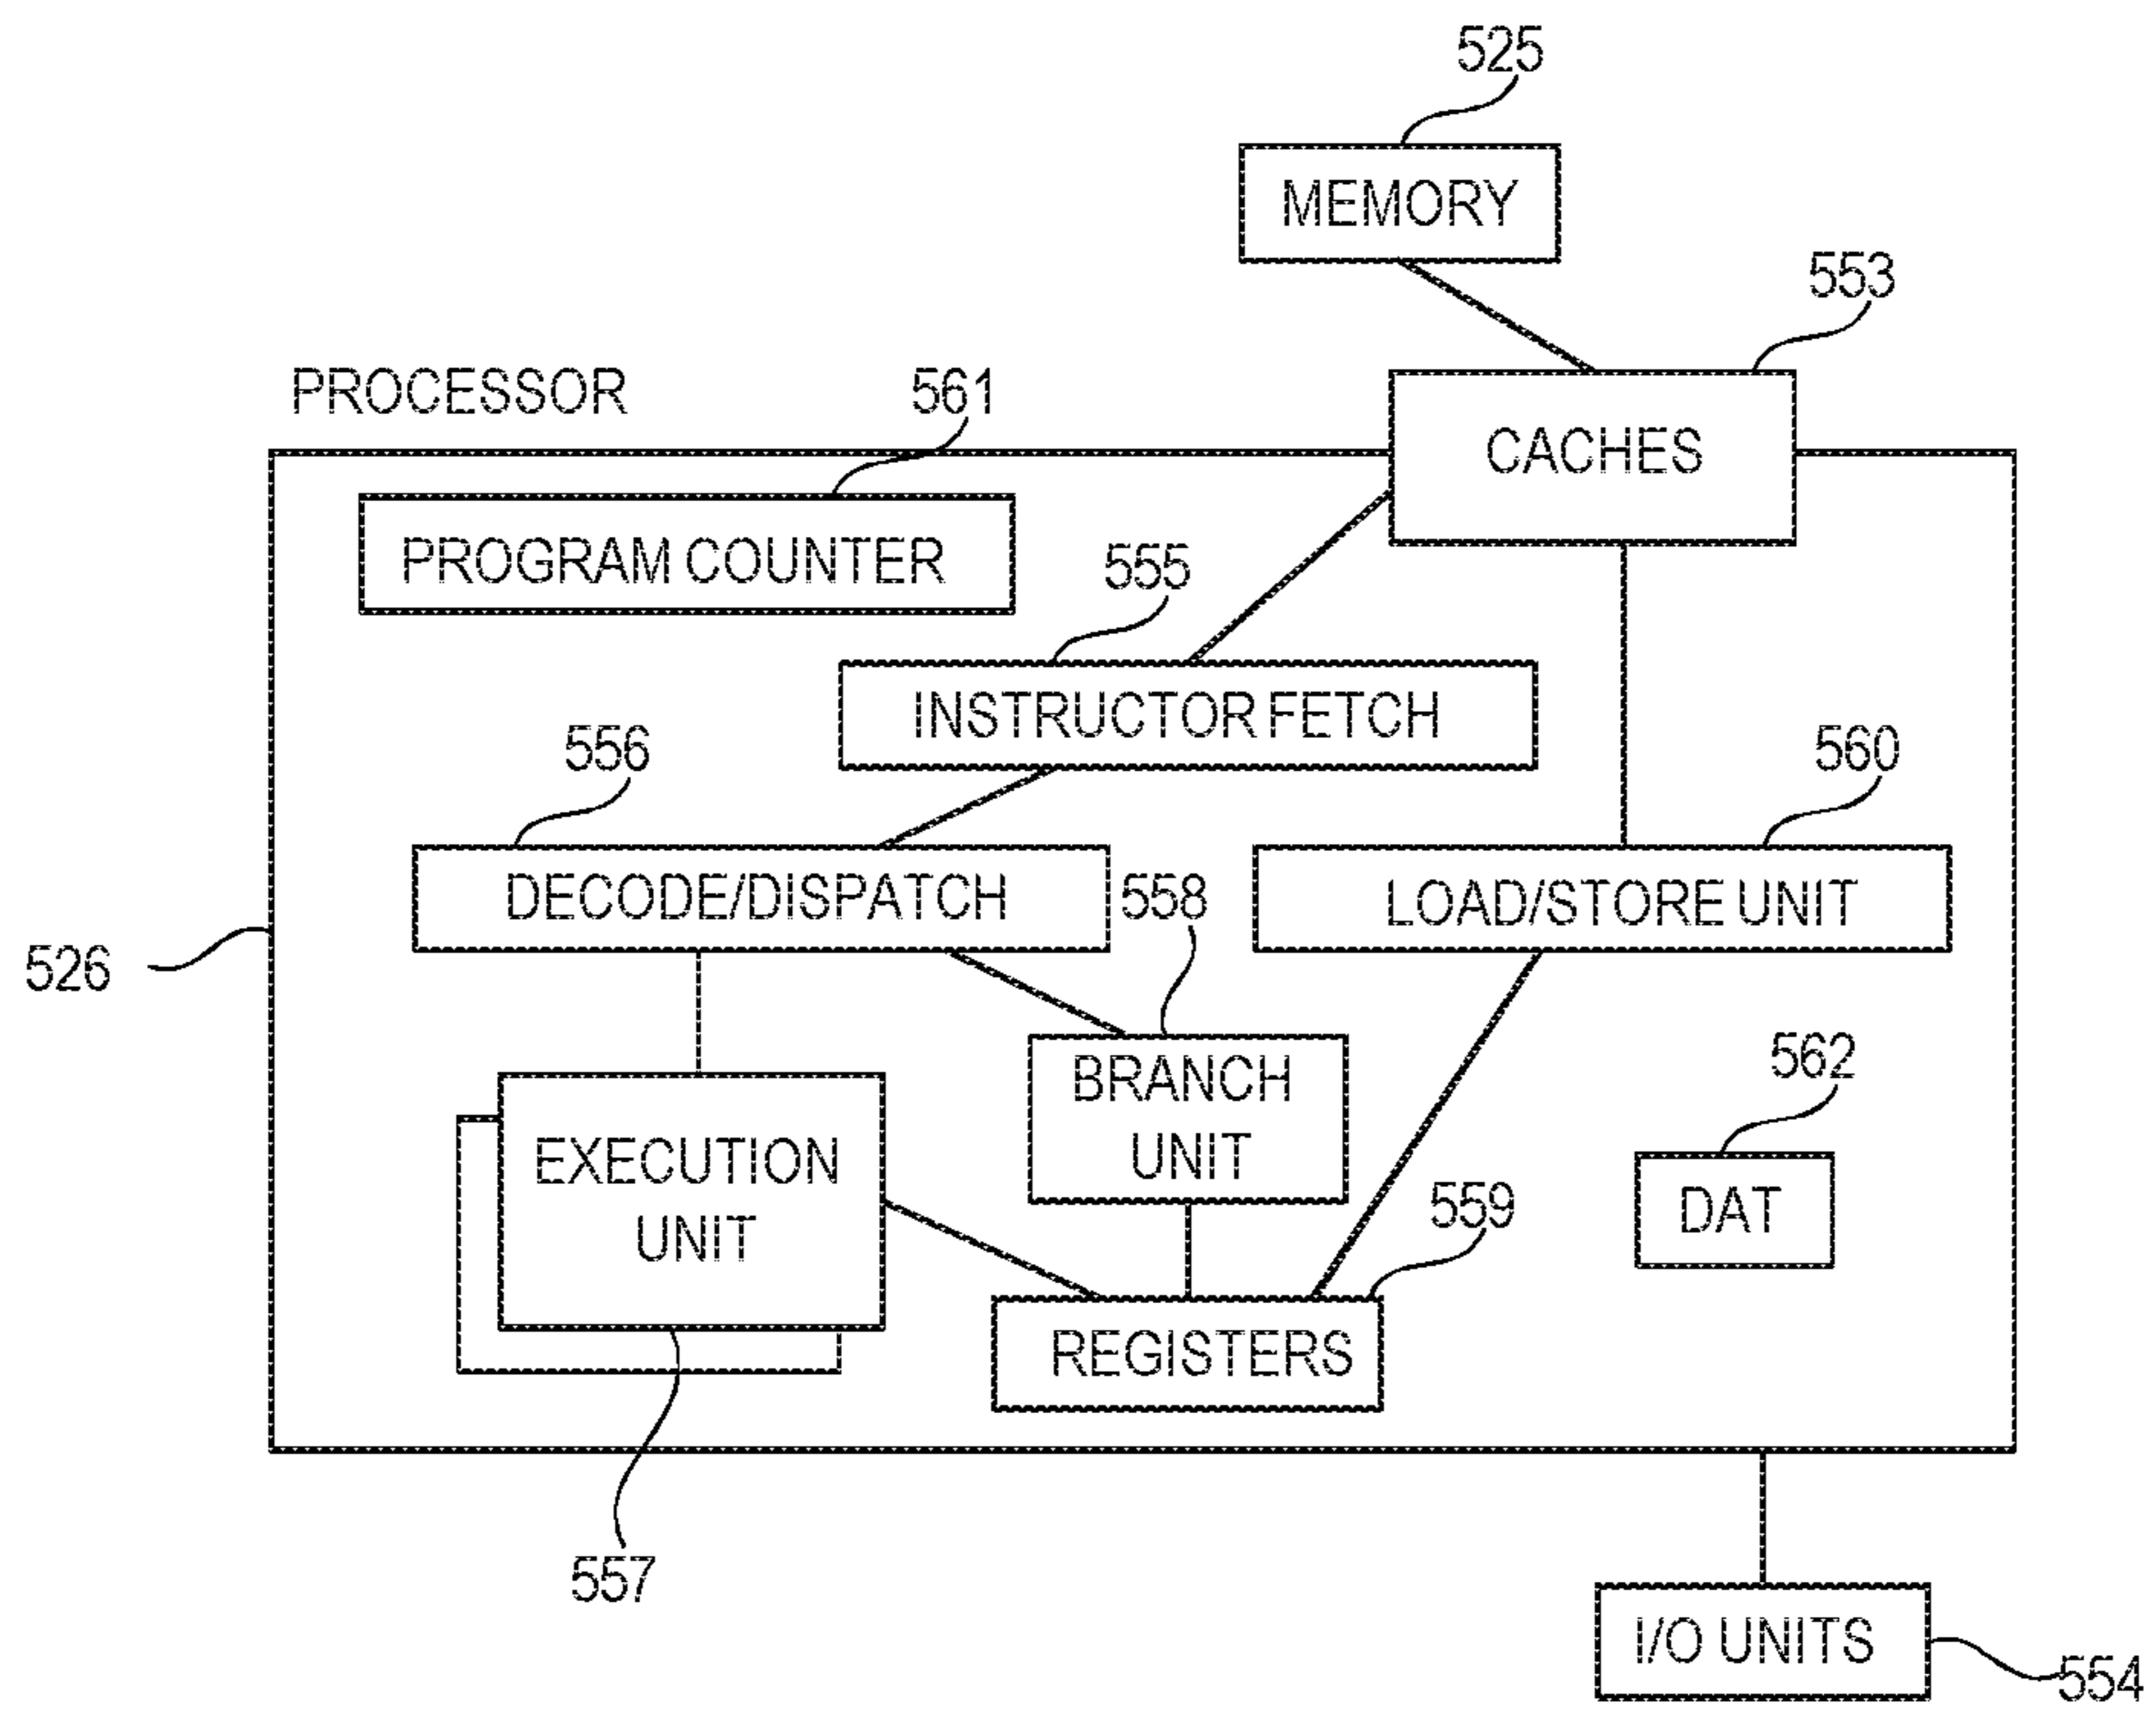


FIG. 21




FIG. 22A




FIG. 22B




FIG. 22C




FIG. 23



**DIRECTED INTERRUPT VIRTUALIZATION**

This application claims priority from European patent application number EP19157093.6, filed Feb. 14, 2019, which is hereby incorporated herein by reference in its entirety.

**BACKGROUND**

One or more aspects of the present disclosure relate, in general, to interrupt processing within a computer system, and in particular, to handling interrupts generated by bus connected modules in a multiprocessor computer system.

Interrupts are used to signal to a processor that an event needs the attention of the processor. For example, hardware devices, e.g., hardware devices connected with the processors via a bus, use interrupts to communicate that they require attention from an operating system. In case the receiving processor is currently performing some activities, the receiving processor may in response to receiving an interrupt signal suspend its current activities, save its state, and handle the interrupt, e.g., by executing an interrupt handler. The interruption of the processors' current activities resulting from the receipt is only temporary. After having handled the interrupt, the processor may resume its suspended activities. Thus, interrupts may allow for a performance improvement by eliminating unproductive waiting time of a processor in polling loops, waiting for external events.

In multiprocessor computer systems, interrupt routing efficiency issues may arise. The challenge is to forward interrupt signals sent by hardware devices, like, e.g., bus connected modules, to a processor of the multiple processors assigned for use by the operating system in an efficient way. This may be particularly challenging in case the interrupt is used to communicate with a guest operating system on a virtual machine. A hypervisor or virtual machine monitor (VMM) creates and runs one or more virtual machines, i.e., guest machines. A virtual machine provides a guest operating system executed on the same machine with a virtual operating platform, while hiding the physical characteristics of the underlying platform. Using multiple virtual machines allows to run multiple operating systems in parallel. Since being executed on a virtual operating platform, the guest operating system's view of the processors may in general differ from an underlying, e.g., physical view of the processors. The guest operating system uses virtual processor IDs to identify processors, which in general do not coincide with underlying logical processor IDs. The hypervisor which manages the execution of the guest operating system defines a mapping between underlying logical processor IDs and virtual processor IDs used by the guest operating system. However, this mapping and the selection of processors scheduled for use by the guest operating system are not static, but may be changed by the hypervisor, while the guest operating system is running, without the knowledge of the guest operating system.

Typically, this challenge is solved by forwarding interrupt signals using broadcasting. When using the broadcasting, the interrupt signal is successively forwarded between the multiple processors until a processor suitable for handling the interrupt signal is met. However, in case of multiple processors the probability that a processor receiving the broadcasted interrupt signal first is indeed suitable for handling the interrupt signal may be rather low. Furthermore,

being suitable for handling the interrupt signal not necessarily means that the respective processor is the best choice for handling the interrupt.

**SUMMARY**

Shortcomings of the prior art are overcome and additional advantages are provided through the provision of a computer program product for providing an interrupt signal to a guest operating system executed using one or more processors of a plurality of processors of a computer system assigned for usage by the guest operating system. The computer program product comprises at least one storage medium readable by at least one processing circuit and storing instructions for execution by the processing circuit for performing a method. The method includes receiving, by a bus attachment device of the computer system from a bus connected module of a plurality of bus connected modules operationally coupled to the plurality of processors via the bus attachment device, an interrupt signal with an interrupt target ID. The interrupt target ID identifies one processor of the plurality of processors assigned for usage by the guest operating system as a target processor for handling the interrupt signal. The bus attachment device translates the interrupt target ID to a logical processor ID of the target processor using a mapping table of the bus attachment device. The mapping table maps interrupt target IDs of the processors assigned for usage by the guest operating system to logical processor IDs of the plurality of processors. The bus attachment device forwards the interrupt signal to the target processor, using the logical processor ID of the target processor to address the target processor directly.

Methods and systems relating to one or more aspects are also described and claimed herein.

Additional features and advantages are realized through the techniques described herein. Other embodiments and aspects are described in detail herein and are considered a part of the claimed aspects.

**BRIEF DESCRIPTION OF THE DRAWINGS**

In the following, embodiments of aspects of the invention are explained in greater detail, by way of example only, making reference to the drawings in which:

FIG. 1 depicts one example of a schematic diagram of an example computer system;

FIG. 2 depicts one example of a schematic diagram of an example virtualization scheme;

FIG. 3 depicts one example of a schematic diagram of an example virtualization scheme;

FIG. 4 depicts one example of a schematic diagram of an example virtualization scheme;

FIG. 5 depicts one example of a schematic diagram of an example computer system;

FIG. 6 depicts one example of a schematic diagram of an example computer system;

FIG. 7 depicts one example of a schematic flow diagram of an example method;

FIG. 8 depicts one example of a schematic flow diagram of an example method;

FIG. 9 depicts one example of a schematic flow diagram of an example method;

FIG. 10 depicts one example of a schematic diagram of an example computer system;

FIG. 11 depicts one example of a schematic flow diagram of an example method;



FIG. 12 depicts one example of a schematic flow diagram of an example method;

FIG. 13 depicts one example of a schematic diagram of an example data structure;

FIG. 14 depicts one example of a schematic diagram of example vector structures;

FIG. 15 depicts one example of a schematic diagram of example vector structures;

FIGS. 16A-16B depict examples of schematic diagrams of example vector structures;

FIGS. 17A-17B depict examples of schematic diagrams of example vector structures;

FIG. 18 depicts one example of a schematic diagram of an example computer system;

FIG. 19 depicts one example of a schematic diagram of an example computer system;

FIG. 20 depicts one example of a schematic diagram of an example computer system;

FIG. 21 depicts one example of a schematic diagram of an example computer system;

FIGS. 22A-22C depict examples of schematic diagrams of example units; and

FIG. 23 depicts one example of a schematic diagram of an example computer system.

#### DETAILED DESCRIPTION

The descriptions of the various embodiments of aspects of the present invention will be presented for purposes of illustration, but are not intended to be exhaustive or limited to the embodiments disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of aspects of the described embodiments. The terminology used herein was chosen to best explain the principles of the embodiments, the practical application or technical improvement over technologies found in the marketplace, or to enable others of ordinary skill in the art to understand the embodiments disclosed herein.

Embodiments may have the beneficial effect of enabling the bus attachment device to address the target processor directly. Thus, the interrupt signal may be targeted by the issuing bus connected module selecting a target processor ID to a particular processor, i.e. target processor, of a multiprocessor computer system. For example, a processor may be selected as a target processor for the interrupt signal, which has performed activities related to the interrupt before. Handling the interrupt signal by the same processor as the respective activities may result in a performance advantage, since in case the same processor is also handling the interrupt signal, the data (e.g., all the data) in context with this interrupt may already be available to the processor and/or stored in a local cache enabling a fast access for the respective processor without requiring extensive cache traffic.

Thus, a broadcasting of the interrupt signal may be avoided for which there is no guarantee that the processor which will handle the interrupt in the end is best suited for this task from a performance point of view, like cache traffic minimization. Rather than presenting the interrupt signal to all processors, each processor trying to handle it and one processor wins, the interrupt signal may be provided directly to a target processor increasing the efficiency of the interrupt signal handling.

An interrupt mechanism may be implemented using directed interrupts. The bus attachment device may be enabled to directly address a target processor using a logical

processor ID of the same processor, when forwarding an interrupt signal for handling to its target processor defined by the issuing bus connected module. Translating interrupt target IDs to logical processor IDs by the bus connected device may further ensure that from a point of view of the guest operating system the same processor is addressed, even though the mapping between interrupt target IDs and logical processor IDs or the selection of processors scheduled for use by the guest operating system may be changed by the hypervisor.

For the translation of the interrupt target ID provided with the interrupt signal to a logical processor ID, a static mapping provided by a mapping table of the bus attachment device may be used. Thus, enabling the bus attachment device to identify and use logical processor IDs to address target processors identified by interrupt requests directly.

In accordance with an embodiment, the interrupt signal is received in a form of a message signaled interrupt comprising the interrupt target ID of the target processor. Using message signaled interrupts (MSI) is a method for a bus connected module, such as a Peripheral Component Interconnect (PCI) or Peripheral Component Interconnect express (PCIe) function, to generate a central processing unit (CPU) interrupt in order to notify a guest operating system using the respective central processing unit of the occurrence of an event or the presence of some status. MSI provides an in-band method of signaling an interrupt, using special in-band messages, thereby avoiding a need for dedicated paths separated from a main data path to send such control information, such as dedicated interruption pins on each device. MSI rather relies on exchanging special messages indicating interrupts through the main data path. When a bus connected module is configured to use MSI, the respective module requests an interrupt by performing an MSI write operation of a specified number of bytes of data to a special address. The combination of this special address, i.e., MSI address, and a unique data value, i.e. MSI data, is termed an MSI vector.

Modern PCIe standard adapters have the capability to present multiple interrupts. MSI-X, for example, permits a bus connected module to allocate up to 2048 interrupts. Thus, targeting individual interrupts to different processors, such as in a high-speed networking application relying on multiprocessor systems, is enabled. MSI-X allows to assign a multitude of interrupts, each with an individual MSI address and MSI data value.

For transmitting an interrupt signal, an MSI-X message may be used. The content of the MSI-X message may be determined using an MSI-X data table. The MSI-X data table local to the bus connected module, i.e. PCIe adapter/function, may be indexed by a number assigned to each interrupt signal, also referred to as an interrupt request (IRQ). The MSI-X data table content is under control of the guest operating system and may be set to the operating system by guidance of hardware and/or firmware. A single PCIe adapter may comprise multiple PCIe functions, each of which may have an independent MSI-X data table. This may, for example, be the case for a single root input/output virtualization (SR-IOV) or multi-function devices.

An interrupt target ID, like, e.g., a virtual processor ID, may be directly encoded as part of a message sent by the bus connected module, like, e.g., an MSI-X message, comprising the interrupt signal. The message, e.g., an MSI-X message, may comprise a requestor ID, i.e. an ID of the bus connected module, the aforementioned interrupt target ID, a DIBV (directed interrupt signal vector) or an AIBV (interrupt signal vector) index, an MSI address and MSI data. An



## 5

MSI-X message may provide 64 bits for the MSI address and 32 bits for the data. A bus connected module may request an interrupt using MSI by performing an MSI write operation of a specific MSI data value to a special MSI address.

The device table is, for instance, a shared table which may be fully indexed by the requestor ID (RID) of the interrupt requestor, i.e., the bus connected module. The bus attachment device remaps and posts the interrupt, i.e., the bus attachment device translates the interrupt target ID and uses the same to directly address the target processor.

A guest operating system may use virtual processor IDs to identify processors in a multiprocessor computer system. Thus, the guest operating system's view of processors may not be identical to a view of an underlying system using logical processor IDs. Bus connected modules providing resources used by a guest operating system may use virtual processor IDs as the resources for communicating with the guest operating system, like, e.g., an MSI-X data table, may be under the control of the guest operating system. As an alternative to a virtual processor ID, any other ID may be defined for the bus connected module to address the processors.

The interrupt is presented to the guest operating system or other software executed thereon, such as other programs, etc. As used herein, the term operating system includes operating system device drivers.

As used herein, the term bus connected module may comprise any type of bus connected module. According to embodiments, the module may be a hardware module, like e.g., a storage function, processing module, network module, cryptographic module, PCI/PCIe adapter, other type of input/output module, etc. According to other embodiments, the module may be a software module, i.e., a function, like, e.g., a storage function, processing function, network function, cryptographic function, PCI/PCIe function, other type of input/output function, etc. Thus, in the examples presented herein, module is used interchangeably with function, e.g., PCI/PCIe function, and adapter, e.g., PCI/PCIe adapter, unless otherwise noted.

Embodiments may have the benefit that an interrupt signal routing mechanism, e.g., MSI-X message routing mechanism, is provided which allows it to keep bus connected modules, e.g., PCIe adapters and functions, as well as device drivers used to operate or control the bus connected modules, unchanged. Furthermore, the hypervisor may be kept from intercepting the underlying architecture for implementing communication between bus connected modules and a guest operating system, e.g., PCIe MSI-X architecture. In other words, changes to the interrupt signal routing mechanism may be implemented outside the hypervisor and bus connected modules.

According to an embodiment, the computer system further comprises a memory, and the bus attachment device is operationally connected with the memory. The method further comprises retrieving by the bus attachment device a copy of a device table entry from a device table stored in the memory. The device table entry comprises a direct signaling indicator indicating whether the target processor is to be addressed directly. If the direct signaling indicator indicates a direct forwarding of the interrupt signal, the forwarding of the interrupt signal using a logical processor ID of the target processor to address the target processor directly is executed, else forwarding by the bus attachment device the interrupt signal for handling to the plurality of processors using broadcasting.

## 6

Embodiments may have the beneficial effect of controlling with the direct signaling indicator whether an interrupt signal is forwarded using direct addressing or broadcasting. Using the direct signaling indicator for each bus connected module an individual predefined selection may be provided whether for interrupt signals received from this bus connected module a direct addressing is to be performed or a broadcasting.

According to an embodiment, the direct signaling indicator is implemented as a single bit. Embodiments may have the beneficial effect that the direct signaling indicator is provided in a form with minimum memory space and is fast and efficiently processable.

According to an embodiment, the direct signaling indicator is set for the guest operating system as a static indicator during initialization of the guest operating system. According to an embodiment, the mapping of the interrupt target IDs of the processors assigned for usage by the guest operating system to logical processor IDs of the plurality of processors is a static mapping defined by the mapping table. Embodiments may have the beneficial effect that the bus attachment device is provided with a mapping enabling the attachment device to perform a direct addressing without extensive data fetching from memory.

According to an embodiment, the bus attachment device checks whether the copy of the device table entry is cached in a local cache operationally connected with the bus attachment device. If the copy of the device table entry is cached, the retrieving of the copy of the device table entry is a retrieving from the respective cache, else the retrieving of the device table entry is a retrieving from the memory. Embodiments may have the beneficial effect to ensure a fast and efficient retrieving of the copy of the device table entry.

According to an embodiment, the memory further comprises an interrupt summary vector, and the device table entry further comprises an interrupt summary vector address indicator indicating a memory address of the interrupt summary vector. The interrupt summary vector comprises an interrupt summary indicator per bus connected module, and each interrupt summary indicator is assigned to a bus connected module indicating whether there is an interrupt signal issued by the respective bus connected module to be handled. The method further comprises using by the bus attachment device the indicated memory address of the interrupt summary vector to update the interrupt summary indicator assigned to the bus connected module from which the interrupt signal is received such that the updated interrupt summary indicator indicates that there is an interrupt signal issued by the respective bus connected module to be handled.

Embodiments may have the beneficial effect of monitoring and recording from which bus connected modules there are interrupt signals to be handled. This information may be particularly helpful in case a broadcast has to be performed, e.g., as a fallback in case a direct addressing fails or is not available. According to an embodiment, the interrupt summary vector is implemented as a contiguous area. Embodiments may have the beneficial effect that the interrupt summary vector is provided in a form with minimum memory space and is fast and efficiently processable. The contiguous area may, e.g., be a single cache line. According to an embodiment, each interrupt summary indicator is implemented as a single bit. Embodiments may have the beneficial effect that the interrupt summary indicators are provided in a form with minimum memory space and is fast and efficiently processable.



According to an embodiment, the memory further comprises a directed interrupt summary vector, and the device table entry further comprises a directed interrupt summary vector address indicator indicating a memory address of the directed interrupt summary vector. The directed interrupt summary vector comprises a directed interrupt summary indicator per interrupt target ID, and each directed interrupt summary indicator is assigned to an interrupt target ID indicating whether there is an interrupt signal addressed to the respective interrupt target ID to be handled. The method further comprises using by the bus attachment device the indicated memory address of the directed interrupt summary vector to update the interrupt summary indicator assigned to the target processor ID to which the received interrupt signal is addressed such that the updated interrupt summary indicator indicates that there is an interrupt signal addressed to the respective interrupt target ID to be handled.

Embodiments may have the beneficial effect of monitoring and recording from which bus connected modules there are interrupt signals to be handled. This information may be particularly helpful in case a broadcast has to be performed, e.g., as a fallback in case a direct addressing fails or is not available. In case the directed interrupt summary indicators are assigned to individual interrupt target IDs, only a single indicator, in one embodiment, may have to be checked in order to determine whether there is an interrupt to be handled by a specific processor.

When an interrupt cannot be delivered directly, e.g., because the hypervisor has not scheduled the target processor, the guest operating system may benefit by delivering the interrupt with the originally intended affinity, i.e. information for which processor the interrupt was intended, using broadcast. In this case, the bus attachment device may set a bit specifying the target processor in the DISB (directed interrupt summary vector) after setting the DIBV (directed interrupt signal vector) and before delivering a broadcast interrupt request to the guest operating system. If a guest operating system receives a broadcast interrupt request, it may therefore identify by scanning and disabling direct interrupt summary indicators in the DISB, e.g., scanning and resetting direct interrupt summary bits, which target processors have an interrupt signal pending as signaled in a DIBV. The guest operating system may thus be enabled to decide whether the interrupt signal is handled by the current processor which received the broadcast or further forwarded to the original target processor.

According to an embodiment, the directed interrupt summary vector is implemented as a contiguous area. Embodiments may have the beneficial effect that the directed interrupt summary vector is provided in a form with minimum memory space and is fast and efficiently processable. The contiguous area may, e.g., be a single cache line. According to an embodiment, each directed interrupt summary indicator is implemented as a single bit. Embodiments may have the beneficial effect that the directed interrupt summary indicators are provided in a form with minimum memory space and is fast and efficiently processable.

According to an embodiment, the memory further comprises one or more interrupt signal vectors, and the device table entry further comprises an interrupt signal vector address indicator indicating a memory address of an interrupt signal vector of the one or more interrupt signal vectors. Each of the interrupt signal vectors comprises one or more signal indicators, and each interrupt signal indicator is assigned to a bus connected module of the one or more bus connected modules and an interrupt target ID indicating whether an interrupt signal has been received from the

respective bus connected module addressed to the respective interrupt target ID. The method further comprises: using by the bus attachment device the indicated memory address of the interrupt signal vector to select the interrupt signal indicator assigned to the bus connected module which issued the received interrupt signal and to the interrupt target ID to which the received interrupt signal is addressed; and updating the selected interrupt signal indicator such that the selected interrupt signal indicator indicates that there is an interrupt signal issued by the respective bus connected module and addressed to the respective interrupt target ID to be handled.

According to an embodiment, each of the interrupt signal vectors comprises an interrupt signal indicator per interrupt target ID assigned to the respective interrupt target ID. Each of the interrupt signal vectors is assigned to an individual bus connected module with the interrupt signal indicators of the respective interrupt signal vector being further assigned to the respective individual bus connected module. Embodiments may have the beneficial effect of enabling the guest operating system to keep track for which target processors a bus connected module has issued interrupt signals to be handled.

According to an embodiment, each of the interrupt signal vectors comprises an interrupt signal indicator per bus connected module assigned to the respective bus connected module. Each of the interrupt signal vectors is assigned to an individual target processor ID with the interrupt signal indicators of the respective interrupt signal vector being further assigned to the respective target processor ID. Embodiments may have the beneficial effect of enabling the guest operating system to keep track from which bus connected modules interrupt signals have been issued to be handled by a specific target processor.

The interrupt signal vectors may be implemented, in one embodiment, as directed interrupt signal vectors ordered depending on the target processor ID, i.e., optimized for taking track of directed interrupts. In other words, an order criterium is the target processor IDs rather than requestor IDs identifying the issuing bus connected modules. Each directed interrupt signal vector may comprise one or more directed interrupt signal indicators depending on the number of bus connected modules.

An ordering of interrupt signal indicators, e.g., in a form of interrupt signaling bits, indicating an individual interrupt signal has been received, e.g., in a form of an MSI-X message, sequentially within a contiguous area of memory, like a cache line, for an individual bus connected module, like, e.g., a PCIe function, may thus be avoided. To enable and/or unenable an interrupt signal indicator, e.g., by setting and/or resetting an interrupt signaling bit, the respective contiguous area of memory is, e.g., to be moved to one of the processors to change the respective interrupt signal indicator accordingly.

It may be intended, in one aspect, that a processor handles, e.g., all indicators for which it is responsible from a guest operating system perspective, i.e., in particular all indicators assigned to the respective processor. This may enable a performance advantage, since in case each processor is handling, e.g., all data assigned to the same a likelihood that data required in this context is provided to the processor and/or stored in a local cache may be high enabling a fast access to the respective data for the processor without requiring extensive cache traffic.

However, each processor trying to handle, e.g., all indicators for which it is responsible may nevertheless lead to a high cache traffic between the processors, as each processor



is to write, e.g., all cache lines for, e.g., all functions, since the indicators assigned to each individual processor may be distributed over, e.g., all contiguous areas, such as cache lines.

The interrupt signaling indicators may be reordered in a form of directed interrupt signaling vectors such that, e.g., all interrupt signaling indicators assigned to the same interrupt target ID are combined in the same contiguous area of memory, e.g., cache line. Thus, a processor intending to handle indicators assigned to the respective processor, i.e., interrupt target ID, may only have to load a single contiguous area of memory. Thus, a contiguous area per interrupt target ID is used rather than a contiguous area per bus connected module. Each processor may only need to scan and update a single contiguous area of memory, e.g., a cache line for, e.g., all interrupt signals received from, e.g., all available bus connected modules targeted to that specific processor as a target processor identified by the interrupt target ID. According to embodiments, an offset may be applied by the hypervisor for a guest operating system to align bits to different offsets.

According to an embodiment, each of the interrupt signal vectors is implemented as a contiguous area in memory. Embodiments may have the beneficial effect of providing the interrupt signal vectors in a form with minimum memory space and is fast and efficiently processable. The contiguous area may, e.g., be a cache line. According to an embodiment, each of the interrupt signal indicators is implemented as a single bit. Embodiments may have the beneficial effect of providing the interrupt signal vectors in a form with minimum memory space and is fast and efficiently processable.

According to an embodiment, the device table entry further comprises a logical partition ID identifying a logical partition to which the guest operating system is assigned, and the forwarding of the interrupt signal by the bus attachment device further comprises forwarding with the interrupt signal the logical partition ID. Embodiments may have the beneficial effect of enabling the receiving processor to check to which guest operating system an interrupt signal is addressed.

According to an embodiment, the bus attachment module comprises a mapping table per logical partition ID. Embodiments may have the beneficial effect of providing, e.g., individual mapping tables per hypervisor and/or guest operating system.

According to an embodiment, the method further comprises retrieving by the bus attachment device an interrupt subclass ID identifying an interrupt subclass to which the received interrupt signal is assigned, and the forwarding of the interrupt signal by the bus attachment device further comprises forwarding with the interrupt signal the interrupt subclass ID.

According to an embodiment, the processors of the computer system are used to execute a plurality of guest operating systems, and the bus attachment device comprises a mapping table for each guest operating system of the plurality of guest operating systems.

According to an embodiment, the method further comprises: receiving by the bus attachment device a request from the bus connected module for a direct memory access for updating status information of the bus connected module in the memory, the status update of the bus connected module triggering the interrupt signal; and upon receiving of the request, performing by the bus attachment device the direct memory access to the memory updating the status information of the bus connected module in the memory.

According to embodiments, the instructions provided on the computer readable non-transitory medium for execution by a processing circuit are configured for performing any of the embodiments of the method for providing an interrupt signal to a guest operating system, as described herein.

According to embodiments, the computer system is further configured to execute any of the embodiments of the method for providing an interrupt signal to a guest operating system, as described herein.

FIG. 1 depicts an example computer system **100** for providing an interrupt signal to a guest operating system. The computer system **100** includes a plurality of processors (e.g., CPUs) **130**, which are used to execute a guest operating system. The computer system **100** furthermore includes a memory **140**, also referred to as storage memory or main memory. Memory **140** may provide memory space, i.e. memory sections, assigned for use by hardware, firmware, and software components comprised by computer system **100**. Memory **140** may be used by hardware and firmware of computer system **100**, as well as by software, e.g., hypervisors, host/guest operating systems, application programs, etc. One or more bus connected modules **120** are operationally connected with the plurality of processors **130**, as well as memory **140**, via a bus **102** and a bus attachment device **110**. The bus attachment device **110** manages the communication between bus connected modules **120** on one hand and processors **130**, as well as memory **140**, on the other hand. Bus connected modules **120** may be connected with the bus **102** directly or via one or more intermediate components, like, e.g., a switch **104**.

The bus connected modules **120** may, e.g., be provided in a form of Peripheral Component Interconnect express (PCIe) modules, also referred to as a PCIe adapter or PCIe functions provided by a PCIe adapter. A PCIe function **120** may issue a request which is sent to a bus attachment device **110**, e.g., a PCI Host Bridge (PHB), also referred to a PCI Bridge Unit (PBU). The bus attachment device **110** receives the requests from the bus connected modules **120**. The requests may, for example, comprise an input/output address that is used to perform a direct memory access (DMA) to memory **140** by the bus attachment device **110** or an input/output address that indicates an interrupt signal, e.g., a message signaled interrupt (MSI).

FIG. 2 depicts an example virtual machine support provided by the computer system **100**. The computer system **100** may include one or more virtual machines (VMs) **202**, as well as at least one hypervisor **200**. A virtual machine support may provide an ability to operate large numbers of virtual machines, each capable of executing a guest operating system **204**, such as z/Linux. Each virtual machine **202** may be capable of functioning as a separate system. Thus, each virtual machine may be independently reset, execute a guest operating system, and run different programs, like application programs. An operating system or application program running in a virtual machine appears to have access to the full and complete computer system. However, in reality only a portion of the available resources of the computers system may be available for usage by the respective operating system or application program.

The virtual machines may use a V=V model, in which the memory assigned to a virtual machine is backed by virtual memory, instead of real memory. Thus, each virtual machine has a virtual linear memory space. The physical resources are owned by hypervisor **200**, such as a VM hypervisor, and the shared physical resources are dispatched by the hypervisor to the guest operating systems, as needed, to meet their processing demands. The V=V virtual machine model



## 11

assumes that the interactions between the guest operating systems and the physical shared machine resources are controlled by the VM hypervisor, since the large number of guests may preclude the hypervisor from simply partitioning and assigning the hardware resources to the configured guest.

The processors **130** are assignable to virtual machines **202** by the hypervisor **200**. A virtual machine **202** may, e.g., be assigned with one or more logical processors. Each of the logical processors may represent all or a share of the physical processors **130** that may be dynamically allocated to the virtual machine **202** by the hypervisor **200**. Virtual machines **202** are managed by a hypervisor **200**. The hypervisor **200** may, for example, be implemented in firmware running on the processors **130** or may be a part of an operating system executed on the computer system **100**. The hypervisor **200** may, for example, be a VM hypervisor, such as z/VM® offered by International Business Machines Corporation, Armonk, N.Y.

FIG. 3 depicts an example multilevel virtual machine support provided by the computer system **100**. In addition to the first level virtualization of FIG. 2, a second level virtualization is provided with a second hypervisor **210** being executed on one of the first level guest operating systems which acts as a host operating system for the second hypervisor **210**. The second hypervisor **210** may manage one or more second level virtual machines **212**, each capable of executing a second level guest operating system **214**.

FIG. 4 depicts an example pattern illustrating the use of different types of IDs to identify processors at different hierarchical levels of the computer system **100**. An underlying firmware **220** may provide logical processor IDs ICPU **222** to identify the processors **130** of the computer system **100**. A first level hypervisor **200** uses the logical processor IDs ICPU **222** to communicate with the processors **130**. The first level hypervisor may provide first virtual processor IDs vCPU **224** for use by guest operating systems **204** or second level hypervisors **210** executed on virtual machines managed by the first level hypervisor **200**. The hypervisor **200** may group the first virtual processor IDs vCPU **224** to provide logical partitions, also referred to as zones, for the guest operating systems **204** and/or hypervisors **210**. First virtual processor IDs vCPU **224** are mapped to the logical processor IDs ICPU **222** by the first level hypervisor **200**. One or more of the first virtual processor IDs vCPU **224** provided by the first level hypervisor **200** may be assigned to each guest operating system **204** or hypervisor **210** executed using the first level hypervisor **200**. The second level hypervisor **210** executed on the first level hypervisor **200** may provide one or more virtual machines executing software, like, e.g., further guest operating systems **214**. For this purpose, the second level hypervisor manages second virtual processor IDs vCPU **226** for use by the second level guest operating system **214** executed on virtual machines of the first level hypervisor **200**. The second virtual processor IDs vCPU **226** are mapped to the first virtual processor IDs vCPU **224** by the second level hypervisor **210**.

A bus connected module **120** addressing processors **130** used by a first/second level guest operating system may use a target processor ID in a form of the first/second virtual processor IDs vCPU **224**, **226** or alternative IDs derived from the first/second virtual processor IDs vCPU **224**, **226**.

FIG. 5 depicts one example of a simplified schematic setup of the computer system **100** illustrating participants in a method of providing interrupt signals to a guest operating system executed on the computer system **100**. The simplified setup for the purpose of illustration includes, for instance, a

## 12

bus connected module (BCM) **120** which sends an interrupt signal to a guest operating system executed on one or more of the processors (CPUs) **130**. The interrupt signal is sent to the bus attachment device **110** with an interrupt target ID (IT\_ID) identifying one of the processors **130** as a target processor. The bus attachment device **110** is an intermediate device managing the communication between the bus connected module **120** and the processors **130**, as well as a memory **140** of the computer system **100**. The bus attachment device **110** receives the interrupt signal and identifies using the interrupt target ID a logical processor ID of the target processor in order to address the respective target processor directly. A directed forwarding to the target processor may improve the efficiency of the data processing, e.g., by reducing cache traffic.

FIG. 6 depicts one example of further details of the computer system **100** of FIG. 5. In one embodiment, the bus attachment device **110** is configured to perform status updates of a status of the bus connected module **120** in a module specific area (MSA) **148** of memory **140**. Such a status update may be executed in response to receiving a direct memory access (DMA) write from the bus connected module **120** specifying the status update to be written into memory **140**.

The memory furthermore includes, for instance, a device table (DT) **144** with a device table entry (DTE) **146** for each bus connected module **120**. Upon receipt of an interrupt signal, e.g., an MSI-X write message with an interrupt target ID identifying a target processor for the interrupt request and a requestor ID identifying an origin of the interrupt request in a form of the bus connected module **120**, the bus attachment device **110** fetches the DTE **146** assigned to the requesting bus connected module **120**. The DTE **146** may indicate whether a directed addressing of the target processor is enabled for the requesting bus connected module **120**, e.g., using a direct signaling indicator (e.g., a dIRQ bit). The bus attachment device updates an entry of a directed interrupt signal vector (DIBV) **162**, as well as a directed interrupt summary vector (DISB) **160** in order to keep track for which of the processors **130** an interrupt signal has been received. The DISB **160** may comprise one entry per interrupt target ID, indicating whether there is an interrupt signal to be handled for this processor **130** from any bus connected module **120**. Each DIBV **162** is assigned to one of the interrupt target IDs, i.e. processors **130**, and may comprise one or more entries. Each entry is assigned to one of the bus connected modules **120**. Thus, the DIBV indicates from which bus connected modules there are interrupt signals for a specific processor **130** to be handled. This may have the advantage that in order to check whether there is any interrupt signal or from which bus connected module **120** there are interrupt signals for a specific processor to be handled. Only a signal entry, e.g., bit, or only a signal vector, e.g. bit vector, has to be read from the memory **140**. According to alternative embodiments, interrupt signal vectors (AIBV) and an interrupt summary vector (AISB) may be used. The AIBVs, as well as the entries of the AISB, are each assigned to a specific bus connected module **120**.

The bus attachment device **110** uses, for instance, a mapping table **112** provided on the bus connected module **110** to translate the interrupt target ID (IT\_ID) to a logical processor ID (ICPU) and forwarding the received interrupt signal to the target processor using the logical processor ID to address the target processor directly. Each processor includes, for instance, firmware, e.g. millicode **132**, to receive and process direct interrupt signals. Firmware may further comprise e.g. microcode and/or macrocode of the



processors **130**. It may comprise hardware-level instructions and/or data structures used in implementation of higher-level machine code. According to embodiments, it may comprise proprietary code that may be delivered as micro-code that includes trusted software or microcode specific to the underlying hardware and controls operating system access to the system hardware.

FIG. 7 is one example of a flowchart of an example method of performing a status update of a bus connected module **120** via the bus attachment device **110** using a DMA write request. In step **300**, a bus connected module may decide to update its status and trigger an interrupt, e.g., in order to indicate a signal completion. In step **310**, the bus connected module initiates a direct memory access (DMA) write via the bus attachment device to a section of the memory assigned to a host running on the computer system, i.e. host memory, in order to update the status of the bus connected module. A DMA is the hardware mechanism that allows peripheral components of a computer system to transfer their I/O data directly to and from main memory without a need to involve system processors. In order to perform a DMA, the bus connected module sends a DMA write request, e.g., in a form of an MSI-X message, to the bus attachment device. In case of PCIe, the bus connected module may, e.g., refer to a PCIe function provided on a PCIe adapter. In step **320**, the bus connected module receives the DMA write request with the status update of the bus connected module and updates the memory using the received update. The update may be executed in an area of the host memory reserved for the respective bus connected module.

FIG. 8 is an example of a flowchart of an example method of providing an interrupt signal to a guest operating system using the computer system **100** of FIG. 6. In step **330**, the bus attachment device receives an interrupt signal, e.g., in a form of an MSI-X write message, sent by the bus connected module. This transmission of the interrupt signal may be executed according to the specifications of the PCI architecture. The MSI-X write message comprises an interrupt target ID identifying a target processor of the interrupt. The interrupt target ID may, e.g., be a virtual processor ID used by the guest operating system to identify processors of the multiprocessor computer system. According to embodiments, the interrupt target ID may be any other ID agreed upon by the guest operating system and the bus connected module in order to be able to identify processors. Such an other ID may, for example, be a result of a mapping of a virtual processor ID. In addition, the MSI-X write message may further comprise an interrupt requestor ID (RID), i.e., the ID of a PCIe function issuing the interrupt request, a vector index defining an offset of a vector entry within a vector, an MSI address, e.g., 64 bit address, as well as MSI data, e.g., 32 bit data. The MSI address and MSI data may indicate that the respective write message is in fact an interrupt request in the form of an MSI message.

In step **340**, the bus attachment device fetches a copy of an entry of a device table stored in the memory. The device table entry (DTE) provides address indicators of one or more vectors or vector entries to be updated in order to indicate that an interrupt signal has been received for the target processor. An address indicator for a vector entry may, e.g., comprise an address of the vector in the memory, as well as an offset within the vector. Furthermore, the DTE may provide a direct signaling indicator indicating whether the target processor is to be addressed directly by the bus attachment device using the interrupt target ID provided with the interrupt signal. Furthermore, the DTE may provide

a logical partition ID, also referred to as a zone ID, and an interrupt subclass ID. The respective copy of the device table entry may be fetched from a cache or from memory. In step **350**, the bus attachment device updates the vectors specified in the DTE.

In step **360**, the bus attachment device checks the direct signaling indicator (e.g., dIRQ bit) provided with the interrupt signal. In case the direct signaling indicator indicates no direct signaling, the bus attachment device forwards the interrupt signal by broadcasting using the zone identifier and the interrupt subclass identifier, in order to provide the interrupt signal to a processor used by the guest operating system. In case the direct signaling indicator indicates no direct signaling, in step **370**, the interrupt signal is forwarded to the processor via broadcasting. The broadcast message comprises the zone ID and/or the interrupt subclass ID. When received by a processor, where the interrupt request is enabled for the zone, a status bit is set atomically, e.g., according to a nest communication protocol. Furthermore, the firmware, e.g., millicode, on this processor interrupts its activities, e.g., program execution, and switches to execute an interrupt handler of the guest operating system. In case the direct signaling indicator indicates direct signaling, in step **380**, the bus attachment device translates the interrupt target ID provided with the interrupt signal to a logical processor ID of a processor assigned for usage by the guest operating system. For the translation, the bus attachment device may use a mapping table comprised by the bus attachment device. The bus attachment device may comprise a mapping table or sub-table per zone, i.e. logical partition.

In step **390**, the bus attachment device forwards the interrupt signal to the target processor using the logical processor ID to address the respective processor directly, i.e., sending a direct message. The direct message may further comprise the zone ID and/or interrupt subclass ID. In step **3100**, firmware, e.g., millicode, of the target processor receives the interrupt. In response, the firmware may interrupt its activities, e.g., program execution, and switches to execute an interrupt handler of the guest operating system. The interrupt may be presented to the guest operating system with a direct signaling indication.

FIG. 9 is one example of an additional flowchart further illustrating a method of FIG. 8, in accordance with an aspect of the present invention. In step **400**, the interrupt message is sent to the bus attachment device. In step **402**, the interrupt message is received. In step **404**, it is checked whether the DTE assigned to the interrupt requestor, i.e., the bus connected module, is cached in a local cache operationally connected with the bus attachment device. In case the DTE is not cached, in step **406**, the respective DTE is fetched from memory by the bus attachment device. Thereafter, or if the DTE is cached, in step **408**, vector address indicators provided by the DTE are used to set the vector bits in the memory. In step **410**, it is checked, using the direct signaling indicator provided by the DTE, whether the target processor is to be addressed directly by the bus attachment device using the interrupt target ID provided with the interrupt signal. In case the target processor is not to be targeted directly, the method continues with broadcasting the interrupt request to the processors in step **412**. In case the target processor is to be targeted directly, the method continues with translating the interrupt target ID to a logical processor ID in step **414** and sending a message forwarding the interrupt signal to the target processor in step **416**. The logical processor ID is used to address the target processor directly. In step **418**, the processor receives the message. In step **420** the processor presents the interrupt request to the



## 15

guest operating system. Then, the processor continues its activities until the next interrupt message is received.

FIG. 10 depicts one example of a further embodiment of the computer system 100. The computer system 100 of FIG. 10 corresponds to, e.g., the computer system 100 of FIG. 6 with the processors 130 in addition comprising check logic 134 to check whether the receiving processor is identical with the target processor according to the interrupt target ID which is forwarded by the bus attachment device 110 to the receiving processor 130. In case the receiving processor 130 is not the target processor, i.e., in case of a mismatch of the received interrupt target ID and a reference interrupt target ID of the receiving processor 130, the interrupt signal is broadcasted to the logical partition in order to find a processor for handling the interrupt signal.

FIG. 11 is one example of a flowchart of an example method of providing an interrupt signal to a guest operating system using, e.g., the computer system 100 of FIG. 10. In addition to the steps shown in FIG. 8, the method according to FIG. 11 comprises the following: In step 390, the bus attachment device forwards the interrupt signal to the target processor using the logical processor ID to address the respective processor directly, i.e., sending a direct message. The direct message comprises the interrupt target IS. The direct message may further comprise the zone ID and/or interrupt subclass ID. The receiving processor comprises an interrupt target ID checking logic. The checking logic may furthermore take a logical partition ID into account, in case the interrupt target ID is unique only per logical partition.

In step 392, the checking logic checks whether the received interrupt target ID and/or logical partition ID match an interrupt target ID and/or logical partition currently assigned to the receiving processor and accessible for the checking logic. In case of a mismatch, the receiving firmware in step 394 initiates a broadcast and broadcasts the received interrupt request to the remaining processors using the logical partition ID and/or an interrupt subclass ID to identify a valid target processor for handling the interrupt. In case of a positive match, the receiving firmware, e.g., millicode, of the target processor accepts the directly addressed interrupt for presentation to the guest operating system in step 3100. In response, the firmware may interrupt its activities, e.g., program execution, and switches to execute an interrupt handler of the guest operating system. The interrupt may be presented to the guest operating system with a direct signaling indication.

FIG. 12 is one example of an additional flowchart further illustrating the method of an aspect of the invention. In addition to the steps shown in FIG. 9, the method according to FIG. 12 comprises the following: The message comprises the interrupt target ID, a logical partition ID and an interrupt subclass ID. In step 418, the processor receives the message. In step 419, the processor checks whether the interrupt target ID and/or logical partition ID match the current interrupt target ID and/or logical partition ID provided as references for the check. In case of a match, the processor presents the interrupt request to the guest operating system in step 420. In case of a mismatch, the processor broadcasts the interrupt request to the other processors in step 422. Then, the processor continues its activities until the next interrupt message is received.

FIG. 13 depicts an example of a DTE 146 comprising, for example, a logical partition ID (zone) and an offset (DIBVO) within a DIBV assigned to the interrupt target ID. The DIBVO identifies the start of a section or entry of the vector assigned to a specific bus connected module. The interrupt signal, e.g. MSI-X message, may provide a DIBV-

## 16

Idx which is added to the DIBVO to identify a specific entry of the vector assigned to the bus connected module. Furthermore, a directed number of interrupts (NOI) is provided defining a maximum number of bits in the DIBV reserved for the respective bus connected module. Further details of the DIBV are illustrated in FIG. 16A. In case of an AIBV, the DTE may provide corresponding AIBV specific parameters as illustrated in FIG. 16B.

FIG. 14 depicts one example of a schematic structure of a DISB 160 and a plurality of DIBVs 162. The DISB 160 may be provided in a form of a contiguous section of memory, e.g., a cache line, comprising an entry 161, e.g. bit, per interrupt target ID. Each entry indicates whether there is an interrupt request (IRQ) to be handled by the respective processor identified by the interrupt target ID. For each interrupt target ID, i.e., entry of the DISB 160, a DIBV 162 is provided. Each DIBV 162 is assigned to a specific interrupt target ID and comprises one or more entries 163 per bus connected module MN A, MN B. The DIBVs 162 may each be provided in a form of a contiguous section of memory, e.g., a cache line, comprising the entries 163 assigned to the same interrupt target ID. Entries of different bus connected modules may be ordered using different offset DIBVOs per bus connected module.

FIG. 15 depicts one example of a schematic structure of an AISB 170 and a plurality of AIBVs 172. The AISB 170 may be provided in a form of a contiguous section of memory, e.g., a cache line, comprising an entry 171, e.g., bit, per bus connected module MN A to MN D. Each entry indicates whether there is an interrupt request (IRQ) to be handled from the respective bus connected module. For each bus connected module, i.e., entry of the AISB 170, an AIBV 172 is provided. Each AIBV 172 is assigned to a specific bus connected module and comprises one or more entries 173 per interrupt target ID. The AIBVs 172 may each be provided in a form of a contiguous section of memory, e.g., a cache line, comprising the entries 173 assigned to the same bus connected module. Entries regarding different target processor IDs may be ordered using different offset AIBVOs per bus connected module.

FIGS. 17A and 17B show an example DISB 160 and AISB 170, respectively. Entries 161, 171 may be addressed using a basic address DISB@ and AISB@, respectively, in combination with an offset DISBO and AISBO, respectively. In case of the DISB 160, DISBO may, e.g., be identical with the interrupt target ID to which the respective entry 161 is assigned. The interrupt target ID may, e.g., be provided in a form of a virtual processor ID (vCPU).

Various embodiments of aspects of the invention provide for a method of providing an interrupt signal to a guest operating system executed using one or more processors of a plurality of processors of a computer system assigned for usage by the guest operating system, as well as a computer system and a computer program product as described by the subject matter of the independent claims. Embodiments are described in the dependent claims. Embodiments of the present invention can be freely combined with each other if they are not mutually exclusive.

One aspect of the invention relates to a method of providing an interrupt signal to a guest operating system executed using one or more processors of a plurality of processors of a computer system assigned for usage by the guest operating system. The computer system further comprises one or more bus connected modules being operationally connected with the plurality of processors via a bus and a bus attachment device. Each processor of the plurality of processors is assigned with a logical processor ID used by



the bus attachment device to address the respective processor, and each processor of the plurality of processors assigned for usage by the guest operating system is further assigned with an interrupt target ID used by the guest operating system and the one or more bus connected modules to address the respective processor. The method comprises receiving by the bus attachment device an interrupt signal from one of the bus connected modules with an interrupt target ID, the interrupt target ID identifying one of the processors assigned for usage by the guest operating system as a target processor for handling the interrupt signal; translating by the bus attachment device the received interrupt target ID to a logical processor ID of the target processor using a mapping table comprised by the bus attachment device, the mapping table mapping the interrupt target IDs of the processors assigned for usage by the guest operating system to logical processor IDs of the plurality of processors; and forwarding by the bus attachment device the interrupt signal to the target processor for handling using the logical processor ID of the target processor to address the target processor directly.

Another aspect of the invention relates to a computer system for providing an interrupt signal to a guest operating system executed using one or more processors of a plurality of processors of the computer system assigned for usage by the guest operating system. The computer system further comprises one or more bus connected modules being operationally connected with the plurality of processors via a bus and a bus attachment device. Each processor of the plurality of processors is assigned with a logical processor ID used by the bus attachment device to address the respective processor, and each processor of the plurality of processors assigned for usage by the guest operating system is further assigned with an interrupt target ID used by the guest operating system and the one or more bus connected modules to address the respective processor. The computer system is configured to perform a method comprising: receiving by the bus attachment device an interrupt signal from one of the bus connected modules with an interrupt target ID, the interrupt target ID identifying one of the processors assigned for usage by the guest operating system as a target processor for handling the interrupt signal; translating by the bus attachment device the received interrupt target ID to a logical processor ID of the target processor using a mapping table comprised by the bus attachment device, the mapping table mapping the interrupt target IDs of the processors assigned for usage by the guest operating system to logical processor IDs of the plurality of processors; and forwarding by the bus attachment device the interrupt signal to the target processor for handling using the logical processor ID of the target processor to address the target processor directly.

Another aspect of the invention relates to a computer program product for providing an interrupt signal to a guest operating system executed using one or more processors of a plurality of processors of a computer system assigned for usage by the guest operating system. The computer system further comprises one or more bus connected modules being operationally connected with the plurality of processors via a bus and a bus attachment device. Each processor of the plurality of processors is assigned with a logical processor ID used by the bus attachment device to address the respective processor, and each processor of the plurality of processors assigned for usage by the guest operating system is further assigned with an interrupt target ID used by the guest operating system and the one or more bus connected modules to address the respective processor. The computer

program product comprises a computer readable non-transitory medium readable by a processing circuit and storing instructions for execution by the processing circuit for performing a method comprising: receiving by the bus attachment device an interrupt signal from one of the bus connected modules with an interrupt target ID, the interrupt target ID identifying one of the processors assigned for usage by the guest operating system as a target processor for handling the interrupt signal; translating by the bus attachment device the received interrupt target ID to a logical processor ID of the target processor using a mapping table comprised by the bus attachment device, the mapping table mapping the interrupt target IDs of the processors assigned for usage by the guest operating system to logical processor IDs of the plurality of processors; and forwarding by the bus attachment device the interrupt signal to the target processor for handling using the logical processor ID of the target processor to address the target processor directly.

In one or more embodiments, the guest operating system may, for example, be implemented using a pageable storage mode guest. A pageable guest, e.g., in a z/Architecture® hardware architecture, may be interpretively executed via a Start Interpretive Execution (SIE) instruction, at level 2 of interpretation. For instance, a logical partition (LPAR) hypervisor executes the SIE instruction to begin a logical partition in physical, fixed memory. An operating system in that logical partition, e.g., z/VM®, may issue the SIE instruction to execute its guests (virtual) machines in its virtual storage. Thus, the LPAR hypervisor may use level-1 SIE and the z/VM® hypervisor may use level-2 SIE.

According to embodiments, a computer system is a System z® server offered by International Business Machines Corporation. System z® is based on the z/Architecture® hardware architecture offered by International Business Machines Corporation. Details regarding the z/Architecture® hardware architecture are described in an IBM publication entitled, “z/Architecture Principles of Operation,” IBM Publication No. SA22-7832-11, Aug. 25, 2017, which is hereby incorporated herein by reference in its entirety. IBM, System z, z/VM, and z/Architecture are registered trademarks or trademarks of International Business Machines Corporation, Armonk, N.Y. in at least one jurisdiction. Other names used herein may be registered trademarks, trademarks or product names of International Business Machines Corporation or other companies.

According to embodiments, computer systems of other architectures may implement and use one or more aspects of the present invention. As examples, servers other than System z® servers, such as Power Systems servers or other servers offered by International Business Machines Corporation, or servers of other companies implement, use and/or benefit from one or more aspects of the present invention. Power is a trademark or registered trademark of International Business Machines Corporation in at least one jurisdiction. Further, although in the examples herein, the bus connected modules and the bus attachment device are considered a part of a server, in other embodiments, they do not have to necessarily be considered a part of the server, but may simply be considered as being coupled to system memory and/or other components of a computer system. The computer system need not be a server. Further, although the bus connected modules may be PCIe, one or more aspects of the present invention are usable with other bus connected modules. PCIe adapter and PCIe functions are just examples. Further, one or more aspects of the present invention may be applicable to interrupt schemes other than PCI MSI and PCI MSI-X. Yet further, although examples are



described in which bits are set, in other embodiments, bytes or other types of indicators may be set. Moreover, the DTE and other structures may include more, less and/or different information.

Further, other types of computer systems may benefit from one or more aspects of the present invention. As an example, a data processing system suitable for storing and/or executing program code is usable that includes at least two processors coupled directly or indirectly to memory elements through a system bus. The memory elements include, for instance, local memory employed during actual execution of the program code, bulk storage, and cache memory which provide temporary storage of at least some program code in order to reduce the number of times code is to be retrieved from bulk storage during execution.

Input/Output or I/O devices including, but not limited to, keyboards, displays, pointing devices, DASD (Direct Access Storage Devices), tape, CDs (Compact Discs), DVDs (Digital Versatile Discs), thumb drives and other memory media, etc., may be coupled to the system either directly or through intervening I/O controllers. Network adapters may also be coupled to the system to enable the data processing system to become coupled to other data processing systems or remote printers or storage devices through intervening private or public networks. Modems, cable modems, and Ethernet cards are just a few of the available types of network adapters.

Referring to FIG. 18, representative components of a host computer system 500 to implement one or more aspects of the present invention are portrayed. The representative host computer 500 comprises one or more processors, e.g., CPUs, 501 in communication with computer memory 502, as well as I/O interfaces to storage media devices 511 and networks 510 for communicating with other computers or SANs and the like. The CPUs 501 are compliant with an architecture having an architected instruction set and architected functionality. The CPUs 501 may have dynamic address translation (DAT) 503 for transforming program addresses, virtual addresses, into real addresses of memory. A DAT may comprise a translation lookaside buffer (TLB) 507 for caching translations so that later accesses to the block of computer memory 502 do not require the delay of address translation. A cache 509 may be employed between computer memory 502 and the CPUs 501. The cache 509 may be hierarchically structured providing a large, high level cache available to more than one CPU and smaller, faster, lower level caches between the high-level cache and each CPU. In some implementations, the lower level caches may be split to provide separate low-level caches for instruction fetching and data accesses. According to embodiments, an instruction may be fetched from memory 502 by an instruction fetch unit 504 via a cache 509. The instruction may be encoded in an instruction decode unit 506 and dispatched, in some embodiments with other instructions, to instruction execution unit or units 508. Several execution units 508 may be employed, for example an arithmetic execution unit, a floating-point execution unit and a branch instruction execution unit. The instruction is executed by the execution unit, accessing operands from instruction specified registers or memory as needed. If an operand is to be accessed, e.g., loaded or stored, from memory 502, a load/store unit 505 may handle the access under control of the instruction being executed. Instructions may be executed in hardware circuits or in internal microcode, i.e., firmware, or by a combination of both.

A computer system may comprise information in local or main storage, as well as addressing, protection, and refer-

ence and change recording. Some aspects of addressing include the format of addresses, the concept of address spaces, the various types of addresses, and the manner in which one type of address is translated to another type of address. Some of main storage includes permanently assigned storage locations. Main storage provides the system with directly addressable fast-access storage of data. Both data and programs are to be loaded into main storage, e.g., from input devices, before they may be processed.

Main storage may include one or more smaller, faster-access buffer storages, sometimes called caches. A cache may be physically associated with a CPU or an I/O processor. The effects, except on performance, of the physical construction and use of distinct storage media may generally not be observable by programs executed.

Separate caches may be maintained for instructions and for data operands. Information within a cache may be maintained in contiguous bytes on an integral boundary called a cache block or cache line. A model may provide an EXTRACT CACHE ATTRIBUTE instruction which returns the size of a cache line in bytes. A model may also provide PREFETCH DATA and PREFETCH DATA RELATIVE LONG instructions which effects the prefetching of storage into the data or instruction cache or the releasing of data from the cache.

Storage may be viewed as a long horizontal string of bits. For most operations, accesses to storage may proceed in a left-to-right sequence. The string of bits is subdivided into units of eight bits. An eight-bit unit is called a byte, which is the basic building block of all information formats. Each byte location in storage may be identified by a unique nonnegative integer, which is the address of that byte location, also referred to as the byte address. Adjacent byte locations may have consecutive addresses, starting with 0 on the left and proceeding in a left-to-right sequence. Addresses are unsigned binary integers and may e.g. be 24, 31, or 64 bits.

Information is transmitted between memory and CPUs one byte, or a group of bytes, at a time. Unless otherwise specified, in, for instance, the z/Architecture® hardware architecture, a group of bytes in memory is addressed by the leftmost byte of the group. The number of bytes in the group is either implied or explicitly specified by the operation to be performed. When used in a CPU operation, a group of bytes is called a field. Within each group of bytes, in, for instance, the z/Architecture® hardware architecture, bits are numbered in a left-to-right sequence. In the z/Architecture® hardware architecture, the leftmost bits are sometimes referred to as the “high-order” bits and the rightmost bits as the “low-order” bits. Bit numbers are not storage addresses, however. Only bytes may be addressable. To operate on individual bits of a byte in storage, the entire byte may be accessed. The bits in a byte may be numbered 0 through 7, from left to right in, e.g., the z/Architecture® hardware architecture. The bits in an address may be numbered 8-31 or 40-63 for 24-bit addresses, or 1-31 or 33-63 for 31-bit addresses; they are numbered 0-63 for 64-bit addresses. Within any other fixed-length format of multiple bytes, the bits making up the format may be consecutively numbered starting from 0. For purposes of error detection, and, e.g., for correction, one or more check bits may be transmitted with each byte or with a group of bytes. Such check bits are generated automatically by the machine and cannot be directly controlled by the program. Storage capacities are expressed in number of bytes. When the length of a storage-operand field is implied by the operation code of an instruction, the field is said to have a fixed length, which may be



one, two, four, eight, or sixteen bytes. Larger fields may be implied for some instructions. When the length of a storage-operand field is not implied but is stated explicitly, the field is said to have a variable length. Variable-length operands may vary in length by increments of one byte or with some instructions, in multiples of two bytes or other multiples. When information is placed in storage, the contents of only those byte locations are replaced that are included in the designated field, even though the width of the physical path to storage may be greater than the length of the field being stored.

Certain units of information are to be on an integral boundary in storage. A boundary is called integral for a unit of information when its storage address is a multiple of the length of the unit in bytes. Special names are given to fields of 2, 4, 8, and 16 bytes on an integral boundary. A halfword is a group of two consecutive bytes on a two-byte boundary and is the basic building block of instructions. A word is a group of four consecutive bytes on a four-byte boundary. A doubleword is a group of eight consecutive bytes on an eight-byte boundary. A quadword is a group of 16 consecutive bytes on a 16-byte boundary. When storage addresses designate halfwords, words, doublewords, and quadwords, the binary representation of the address contains one, two, three, or four rightmost zero bits, respectively. Instructions are to be on two-byte integral boundaries. The storage operands of most instructions do not have boundary-alignment requirements.

On devices that implement separate caches for instructions and data operands, a significant delay may be experienced if the program stores into a cache line from which instructions are subsequently fetched, regardless of whether the store alters the instructions that are subsequently fetched.

In one embodiment, the aspects of the invention may be practiced by software, sometimes referred to licensed internal code, firmware, micro-code, milli-code, pico-code and the like, any of which would be consistent with aspects of the present invention. Referring to FIG. 18, software program code which embodies aspects of the present invention may be accessed from long-term storage media devices 511, such as a CD-ROM drive, tape drive or hard drive. The software program code may be embodied on any of a variety of known media for use with a data processing system, such as a diskette, hard drive, or CD-ROM. The code may be distributed on such media, or may be distributed to users from computer memory 502 or storage of one computer system over a network 510 to other computer systems for use by users of such other systems.

Software program code may comprise an operating system which controls the function and interaction of the various computer components and one or more application programs. Program code may be paged from storage media device 511 to the relatively higher-speed computer storage 502 where it is available for processing by processor 501. Well known techniques and methods for embodying software program code in memory, on physical media, and/or distributing software code via networks may be used. Program code, when created and stored on a tangible medium, including but not limited to electronic memory modules (RAM), flash memory, Compact Discs (CDs), DVDs, Magnetic Tape, may be referred to as a "computer program product". The computer program product medium may be readable by a processing circuit in, e.g., a computer system for execution by the processing circuit.

FIG. 19 illustrates a representative workstation or server hardware system in which embodiments of aspects of the present invention may be implemented. The system 520 of

FIG. 19 comprises, e.g., a representative base computer system 521, such as a personal computer, a workstation or a server, including optional peripheral devices. The base computer system 521 includes one or more processors 526 and a bus employed to connect and enable communication between the processor(s) 526 and the other components of the system 521 in accordance with known techniques. The bus connects the processor 526 to memory 525 and long-term storage 527 which may include a hard drive, including any of magnetic media, CD, DVD and Flash Memory for example, or a tape drive for example. The system 521 might also include a user interface adapter, which connects the microprocessor 526 via the bus to one or more interface devices, such as a keyboard 524, a mouse 523, a printer/scanner 530 and/or other interface devices, which may be any user interface device, such as a touch sensitive screen, digitized entry pad, etc. The bus also connects a display device 522, such as an LCD screen or monitor, to the microprocessor 526 via a display adapter.

The system 521 may communicate with other computers or networks of computers by way of a network adapter capable of communicating 528 with a network 529. Example network adapters are communications channels, token ring, Ethernet or modems. Alternatively, the system 521 may communicate using a wireless interface, such as a cellular digital packet data (CDPD) card. The system 521 may be associated with such other computers in a Local Area Network (LAN) or a Wide Area Network (WAN), or the system 521 may be a client in a client/server arrangement with another computer, etc.

FIG. 20 illustrates a data processing network 540 in which embodiments of the present invention may be implemented. The data processing network 540 may include a plurality of individual networks, such as a wireless network and a wired network, each of which may include a plurality of individual workstations 541, 542, 543, 544. Additionally, as those skilled in the art will appreciate, one or more LANs may be included, where a LAN may comprise a plurality of intelligent workstations coupled to a host processor.

Still referring to FIG. 20, the networks may also include mainframe computers or servers, such as a gateway computer, e.g., client server 546, or application server, e.g., remote server 548 which may access a data repository and may also be accessed directly from a workstation 545. A gateway computer 546 may serve as a point of entry into each individual network. A gateway may be needed when connecting one networking protocol to another. The gateway 546 may be, e.g., coupled to another network, like the Internet 547 for example, by means of a communications link. The gateway 546 may also be directly coupled to one or more workstations 541, 542, 543, 544 using a communications link. The gateway computer may be implemented utilizing an IBM eServer™ System z® server available from International Business Machines Corporation.

Referring concurrently to FIG. 19 and FIG. 20, software programming code which may embody aspects of the present invention may be accessed by the processor 526 of the system 520 from long-term storage media 527, such as a CD-ROM drive or hard drive. The software programming code may be embodied on any of a variety of known media for use with a data processing system, such as a diskette, hard drive, or CD-ROM. The code may be distributed on such media, or may be distributed to users 550, 551 from the memory or storage of one computer system over a network to other computer systems for use by users of such other systems.



Alternatively, the programming code may be embodied in the memory **525**, and accessed by the processor **526** using the processor bus. Such programming code may include an operating system **531** which controls the function and interaction of the various computer components and one or more application programs **532**. Program code may be paged from storage media **527** to high-speed memory **525**, where it is available for processing by the processor **526**. Well known techniques and methods for embodying software programming code in memory, on physical media, and/or distributing software code via networks may be used.

The cache that is most readily available to a processor, i.e., which may be faster and smaller than other caches of the processor, is the lowest cache, also referred to as the L1 or level one cache, and main memory is the highest-level cache, also referred to as Ln, e.g., L3, if there are n, e.g., n=3, levels. The lowest-level cache may be divided into an instruction cache, also referred to as I-cache, holding machine readable instructions to be executed and a data cache, also referred to as D-Cache, holding data operands.

Referring to FIG. **21**, an example processor embodiment is depicted for processor **526**. One or more levels of cache **553** may be employed to buffer memory blocks in order to improve processor performance. The cache **553** is a high-speed buffer holding cache lines of memory data that are likely to be used. Cache lines may for example be 64, 128 or 256 bytes of memory data. Separate caches may be employed for caching instructions and for caching data. Cache coherence, i.e., synchronization of copies of lines in memory and the caches, may be provided by various suitable algorithms, e.g., "snoop" algorithms. Main memory storage **525** of a processor system may be referred to as a cache. In a processor system having 4 levels of cache **553**, main storage **525** is sometimes referred to as the level 5 (L5) cache, since it may be faster and only holds a portion of the non-volatile storage that is available to a computer system. Main storage **525** "caches" pages of data paged in and out of the main storage **525** by the operating system.

A program counter (instruction counter) **561** keeps track of the address of the current instruction to be executed. A program counter in a z/Architecture® processor is 64 bits and may be truncated to 31 or 24 bits to support prior addressing limits. A program counter may be embodied in a program status word (PSW) of a computer such that it persists during context switching. Thus, a program in progress, having a program counter value, may be interrupted by, for example, the operating system resulting in a context switch from the program environment to the operating system environment. The PSW of the program maintains the program counter value while the program is not active, and the program counter in the PSW of the operating system is used while the operating system is executing. The program counter may be incremented by an amount equal to the number of bytes of the current instruction. Reduced Instruction Set Computing (RISC) instructions may be fixed length, while Complex Instruction Set Computing (CISC) instructions may be variable length. Instructions of the IBM z/Architecture® hardware architecture are, e.g., CISC instructions having a length of 2, 4 or 6 bytes. The Program counter **561** may be modified by either a context switch operation or a branch taken operation of a branch instruction for example. In a context switch operation, the current program counter value is saved in the program status word along with other state information about the program being executed, such as condition codes, and a new program counter value is loaded pointing to an instruction of a new program module to be executed. A branch taken operation

may be performed in order to permit the program to make decisions or loop within the program by loading the result of the branch instruction into the program counter **561**.

An instruction fetch unit **555** may be employed to fetch instructions on behalf of the processor **526**. The fetch unit either fetches "next sequential instructions," target instructions of branch taken instructions, or first instructions of a program following a context switch. Modern instruction fetch units may employ prefetch techniques to speculatively prefetch instructions based on the likelihood that the prefetched instructions might be used. For example, a fetch unit may fetch 16 bytes of an instruction that includes the next sequential instruction and additional bytes of further sequential instructions.

The fetched instructions may then be executed by the processor **526**. According to embodiments, the fetched instruction(s) may be passed to a dispatch unit **556** of the fetch unit. The dispatch unit decodes the instruction(s) and forwards information about the decoded instruction(s) to appropriate units **557**, **558**, **560**. An execution unit **557** may receive information about decoded arithmetic instructions from the instruction fetch unit **555** and may perform arithmetic operations on operands according to the opcode of the instruction. Operands may be provided to the execution unit **557** either from memory **525**, architected registers **559** or from an immediate field of the instruction being executed. Results of the execution, when stored, may be stored either in memory **525**, registers **559** or in other machine hardware, such as control registers, PSW registers and the like.

A processor **526** may have dynamic address translation (DAT) **562** for translating program addresses, e.g., virtual addresses, to, e.g., real addresses, and processor **526** may comprise one or more units **557**, **558**, **560** for executing the function of the instruction. Referring to FIG. **22A**, an execution unit **557** may communicate with architected general registers **559**, a decode/dispatch unit **556**, a load store unit **560**, and other **565** processor units by way of interfacing logic **571**. An execution unit **557** may employ several register circuits **567**, **568**, **569** to hold information that the arithmetic logic unit (ALU) **566** operates on. The ALU performs arithmetic operations such as add, subtract, multiply and divide, as well as logical function such as And, Or, Exclusive-or (XOR), Rotate and Shift. The ALU may support specialized operations that are design dependent. Other circuits may provide other architected facilities **572** including condition codes and recovery support logic for example. The result of an ALU operation may be held in an output register circuit **570** which is configured to forward the result to a variety of other processing functions. There are many arrangements of processor units, the present description is only intended to provide a representative understanding of one embodiment.

An ADD instruction for example may be executed in an execution unit **557** having arithmetic and logical functionality while a floating-point instruction for example would be executed in a floating-point execution having specialized floating point capability. An execution unit operates on operands identified by an instruction by performing an opcode defined function on the operands. For example, an ADD instruction may be executed by an execution unit **557** on operands found in two registers **559** identified by register fields of the instruction.

The execution unit **557** performs the arithmetic addition on two operands and stores the result in a third operand where the third operand may be a third register or one of the two source registers. The execution unit utilizes, e.g., an Arithmetic Logic Unit (ALU) **566** that is capable of per-



forming a variety of logical functions such as Shift, Rotate, And, Or and XOR, as well as a variety of algebraic functions including any of add, subtract, multiply, divide. Some ALUs **566** are designed for scalar operations and some for floating point. Data may be big endian, where the least significant byte is at the highest byte address, or little endian, where the least significant byte is at the lowest byte address, depending on architecture. The IBM z/Architecture® hardware architecture is big endian. Signed fields may be sign and magnitude, 1's complement or 2's complement depending on architecture. A 2's complement number may be advantageous in that the ALU does not need to design a subtract capability since either a negative value or a positive value in 2's complement requires only an addition within the ALU. Numbers may be described in shorthand, where a 12-bit field defines an address of a 4,096-byte block and described as a 4 Kbyte (Kilo-byte) block, for example.

Referring to FIG. 22B, branch instruction information for executing a branch instruction may be sent to a branch unit **558** which often employs a branch prediction algorithm such as a branch history table **582** to predict the outcome of the branch before other conditional operations are complete. The target of the current branch instruction will be fetched and speculatively executed before the conditional operations are complete. When the conditional operations are completed the speculatively executed branch instructions are either completed or discarded based on the conditions of the conditional operation and the speculated outcome. A branch instruction may test condition codes and branch to a target address if the condition codes meet the branch requirement of the branch instruction, a target address may be calculated based on several numbers including ones found in register fields or an immediate field of the instruction for example. The branch unit **558** may employ an ALU **574** having a plurality of input register circuits **575**, **576**, **577** and an output register circuit **580**. The branch unit **558** may communicate with general registers **559**, decode dispatch unit **556** or other circuits **573**, for example, by way of, e.g., interfacing logic **581**.

The execution of a group of instructions may be interrupted for a variety of reasons including a context switch initiated by an operating system, a program exception or error causing a context switch, an I/O interrupt signal causing a context switch or a multi-threading activity of a plurality of programs in a multi-threaded environment, for example. A context switch action saves state information about a currently executing program and then loads state information about another program being invoked. State information may be saved in hardware registers or in memory, for example. State information comprises, e.g., a program counter value pointing to a next instruction to be executed, condition codes, memory translation information and architected register content. A context switch activity may be exercised by hardware circuits, application programs, operating system programs or firmware code, like e.g. microcode, pico-code or licensed internal code (LIC), alone or in combination.

A processor accesses operands according to instruction defined methods. The instruction may provide an immediate operand using the value of a portion of the instruction, may provide one or more register fields explicitly pointing to either general purpose registers or special purpose registers, like e.g. floating-point registers, for example. The instruction may utilize implied registers identified by an opcode field as operands. The instruction may utilize memory locations for operands. A memory location of an operand may be provided by a register, an immediate field, or a

combination of registers and immediate field as exemplified by the z/Architecture® hardware architecture long displacement facility wherein the instruction defines a base register, an index register and an immediate field, i.e., displacement field, that are added together to provide the address of the operand in memory for example. Location herein may imply a location in main memory unless otherwise indicated.

Referring to FIG. 22C, a processor accesses storage using a load/store unit **560**. The load/store unit **560** may perform a load operation by obtaining the address of the target operand in memory **553** and loading the operand in a register **559** or another memory **553** location, or may perform a store operation by obtaining the address of the target operand in memory **553** and storing data obtained from a register **559** or another memory **553** location in the target operand location in memory **553**. The load/store unit **560** may be speculative and may access memory in a sequence that is out-of-order relative to instruction sequence, however the load/store unit **560** is to maintain the appearance to programs that instructions were executed in order. A load/store unit **560** may communicate with general registers **559**, decode/dispatch unit **556**, cache/memory interface **553** or other elements **583** via, e.g., interfacing logic **584**, and comprises various register circuits **586**, **587**, **588** and **589**, ALUs **585** and control logic **590** to calculate storage addresses and to provide pipeline sequencing to keep operations in-order. Some operations may be out-of-order but the load/store unit provides functionality to make the out-of-order operations to appear to the program as having been performed in order.

Addresses that an application program “sees” are often referred to as virtual addresses. Virtual addresses are sometimes also referred to as “logical addresses” and “effective addresses”. These virtual addresses are virtual in that they are redirected to physical memory location by one of a variety of dynamic address translation (DAT) technologies including, but not limited to, simply prefixing a virtual address with an offset value, translating the virtual address via one or more translation tables, the translation tables comprising at least a segment table and a page table alone or in combination, the segment table having an entry pointing to the page table. In the z/Architecture® hardware architecture, a hierarchy of translation is provided including a region first table, a region second table, a region third table, a segment table and an optional page table. The performance of the address translation is often improved by utilizing a translation lookaside buffer (TLB) which comprises entries mapping a virtual address to an associated physical memory location. The entries are created when the DAT translates a virtual address using the translation tables. Subsequent use of the virtual address may then utilize the entry of the fast TLB rather than the slow sequential translation table accesses. TLB content may be managed by a variety of replacement algorithms including Least Recently used (LRU).

Each processor of a multiprocessor system has the responsibility to keep shared resources, such as I/O, caches, TLBs and memory, interlocked for coherency. So-called “snoop” technologies may be utilized in maintaining cache coherency. In a snoop environment, each cache line may be marked as being in any one of a shared state, an exclusive state, a changed state, an invalid state and the like in order to facilitate sharing.

I/O units **554** (FIG. 21) may provide the processor with means for attaching to peripheral devices including tape, disc, printers, displays, and networks for example. I/O units are often presented to the computer program by software



drivers. In mainframes, such as the System z® system from IBM®, channel adapters and open system adapters are I/O units of the mainframe that provide the communications between the operating system and peripheral devices.

Further, other types of computer systems may benefit from one or more aspects of the present invention. As an example, the computer system may comprise an emulator, e.g., software or other emulation mechanisms, in which a particular architecture including, for example, instruction execution, architected functions, such as address translation, and architected registers, is emulated or in which a subset thereof is emulated, e.g., on a native computer system having a processor and memory. In such an environment, one or more emulation functions of the emulator may implement one or more aspects of the present invention, even though a computer executing the emulator may have a different architecture than the capabilities being emulated. For example, in emulation mode, the specific instruction or operation being emulated may be decoded, and an appropriate emulation function may be built to implement the individual instruction or operation.

In an emulation environment, a host computer may for example comprise a memory to store instructions and data, an instruction fetch unit to fetch instructions from memory and to optionally provide local buffering for the fetched instruction, an instruction decode unit to receive the fetched instructions and to determine the type of instructions that have been fetched, and an instruction execution unit to execute the instructions. Execution may include loading data into a register from memory, storing data back to memory from a register, and/or performing some type of arithmetic or logical operation, as determined by the decode unit. For example, each unit may be implemented in software. Operations being performed by the units may be implemented as one or more subroutines within emulator software.

More particularly, in a mainframe, architected machine instructions are used by programmers, like e.g. “C” programmers, e.g., by way of a compiler application. These instructions stored in the storage medium may be executed natively in a z/Architecture® hardware architecture IBM® Server, or alternatively in machines executing other architectures. They may be emulated in the existing and in future IBM® mainframe servers and on other machines of IBM® (e.g., Power Systems servers and System x® Servers). They may be executed in machines running Linux on a wide variety of machines using hardware manufactured by IBM®, Intel®, AMD™, and others. Besides execution on that hardware under a z/Architecture® hardware architecture, Linux may be used as well as machines which use emulation by Hercules, UMX, or FSI (Fundamental Software, Inc), where generally execution is in an emulation mode. In emulation mode, emulation software is executed by a native processor to emulate the architecture of an emulated processor.

A native processor may execute emulation software comprising either firmware or a native operating system to perform emulation of the emulated processor. The emulation software is responsible for fetching and executing instructions of the emulated processor architecture. The emulation software maintains an emulated program counter to keep track of instruction boundaries. The emulation software may fetch one or more emulated machine instructions at a time and convert the one or more emulated machine instructions to a corresponding group of native machine instructions for execution by the native processor. These converted instructions may be cached such that a faster conversion may be accomplished. Notwithstanding, the emulation software is to

maintain the architecture rules of the emulated processor architecture so as to assure operating systems and applications written for the emulated processor operate correctly. Furthermore, the emulation software is to provide resources identified by the emulated processor architecture including, but not limited to, control registers, general purpose registers, floating point registers, dynamic address translation function including segment tables and page tables for example, interrupt mechanisms, context switch mechanisms, Time of Day (TOD) clocks and architected interfaces to I/O subsystems such that an operating system or an application program designed to run on the emulated processor, may be run on the native processor having the emulation software.

A specific instruction being emulated is decoded, and a subroutine is called to perform the function of the individual instruction. An emulation software function emulating a function of an emulated processor is implemented, for example, in a “C” subroutine or driver, or some other method of providing a driver for the specific hardware.

In FIG. 23, an example of an emulated host computer system 592 is provided that emulates a host computer system 500' of a host architecture. In the emulated host computer system 592, the host processor, i.e., CPU 591 is an emulated host processor or virtual host processor and comprises an emulation processor 593 having a different native instruction set architecture than that of the processor 591 of the host computer 500'. The emulated host computer system 592 has memory 594 accessible to the emulation processor 593. In the example embodiment, the memory 594 is partitioned into a host computer memory 596 portion and an emulation routines 597 portion. The host computer memory 596 is available to programs of the emulated host computer 592 according to host computer architecture. The emulation processor 593 executes native instructions of an architected instruction set of an architecture other than that of the emulated processor 591, the native instructions obtained from emulation routines memory 597, and may access a host instruction for execution from a program in host computer memory 596 by employing one or more instructions obtained in a sequence and access/decode routine which may decode the host instructions accessed to determine a native instruction execution routine for emulating the function of the host instruction accessed. Other facilities that are defined for the host computer system 500' architecture may be emulated by architected facilities routines, including such facilities as general purpose registers, control registers, dynamic address translation and I/O subsystem support and processor cache, for example. The emulation routines may also take advantage of functions available in the emulation processor 593, such as general registers and dynamic translation of virtual addresses, to improve performance of the emulation routines. Special hardware and off-load engines may also be provided to assist the processor 593 in emulating the function of the host computer 500'.

It is understood that one or more of the aforementioned embodiments of the invention may be combined as long as the combined embodiments are not mutually exclusive. Ordinal numbers, like e.g. ‘first’ and ‘second’, are used herein to indicate different elements assigned with the same name, but do not necessarily establish any order of the respective elements.

Aspects of the present invention are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations



and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.

Aspects of the present invention may be a system, a method, and/or a computer program product at any possible technical detail level of integration. The computer program product may include a computer readable storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out aspects of the present invention.

The computer readable storage medium can be a tangible device that can retain and store instructions for use by an instruction execution device. The computer readable storage medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable combination of the foregoing. A non-exhaustive list of more specific examples of the computer readable storage medium includes the following: a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a mechanically encoded device such as punch-cards or raised structures in a groove having instructions recorded thereon, and any suitable combination of the foregoing. A computer readable storage medium, as used herein, is not to be construed as being transitory signals per se, such as radio waves or other freely propagating electromagnetic waves, electromagnetic waves propagating through a waveguide or other transmission media (e.g., light pulses passing through a fiber-optic cable), or electrical signals transmitted through a wire.

Computer readable program instructions described herein can be downloaded to respective computing/processing devices from a computer readable storage medium or to an external computer or external storage device via a network, for example, the Internet, a local area network, a wide area network and/or a wireless network. The network may comprise copper transmission cables, optical transmission fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers. A network adapter card or network interface in each computing/processing device receives computer readable program instructions from the network and forwards the computer readable program instructions for storage in a computer readable storage medium within the respective computing/processing device.

Computer readable program instructions for carrying out operations of the present invention may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, configuration data for integrated circuitry, or either source code or object code written in any combination of one or more programming languages, including an object oriented programming language such as Smalltalk, C++, or the like, and procedural programming languages, such as the "C" programming language or similar programming languages. The computer readable program instructions may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN)

or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider). In some embodiments, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays (FPGA), or programmable logic arrays (PLA) may execute the computer readable program instructions by utilizing state information of the computer readable program instructions to personalize the electronic circuitry, in order to perform aspects of the present invention.

Aspects of the present invention are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.

These computer readable program instructions may be provided to a processor of a computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These computer readable program instructions may also be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a particular manner, such that the computer readable storage medium having instructions stored therein comprises an article of manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram block or blocks.

The computer readable program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other device to cause a series of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the instructions which execute on the computer, other programmable apparatus, or other device implement the functions/acts specified in the flowchart and/or block diagram block or blocks.

The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of instructions, which comprises one or more executable instructions for implementing the specified logical function(s). In some alternative implementations, the functions noted in the blocks may occur out of the order noted in the Figures. For example, two blocks shown in succession may, in fact, be accomplished as one step, executed concurrently, substantially concurrently, in a partially or wholly temporally overlapping manner, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts or carry out combinations of special purpose hardware and computer instructions.



Possible combinations of features described herein may be the following:

In one aspect, a method for providing an interrupt signal to a guest operating system executed using one or more processors of a plurality of processors of a computer system assigned for usage by the guest operating system is provided. The computer system further comprises one or more bus connected modules operationally connected with the plurality of processors via a bus and a bus attachment device. Each processor of the plurality of processors is assigned with a logical processor ID used by the bus attachment device to address the respective processor, and each processor of the plurality of processors assigned for usage by the guest operating system is further assigned with an interrupt target ID used by the guest operating system and the one or more bus connected modules to address the respective processor. The method comprises: receiving by the bus attachment device an interrupt signal from one of the bus connected modules with an interrupt target ID, the interrupt target ID identifying one of the processors assigned for usage by the guest operating system as a target processor for handling the interrupt signal; translating by the bus attachment device the received interrupt target ID to a logical processor ID of the target processor using a mapping table comprised by the bus attachment device, the mapping table mapping the interrupt target IDs of the processors assigned for usage by the guest operating system to logical processor IDs of the plurality of processors; and forwarding by the bus attachment device the interrupt signal to the target processor for handling using the logical processor ID of the target processor to address the target processor directly.

In one embodiment, the interrupt signal is received in a form of a message signaled interrupt comprising the interrupt target ID of the target processor.

In one embodiment, the computer system further comprises a memory, and the bus attachment device is operationally connected with the memory. The method further comprises: retrieving by the bus attachment device a copy of a device table entry from a device table stored in the memory, the device table entry comprising a direct signaling indicator indicating whether the target processor is to be addressed directly; if the direct signaling indicator indicates a direct forwarding of the interrupt signal, the forwarding of the interrupt signal using the logical processor ID of the target processor to address the target processor directly is executed, else forwarding by the bus attachment device the interrupt signal for handling to the plurality of processors using broadcasting.

In one embodiment, the direct signaling indicator is implemented as a single bit.

Further, in one embodiment, the direct signaling indicator is set for the guest operating system as a static indicator during initialization of the guest operating system.

In one embodiment, the mapping of the interrupt target IDs of the processors assigned for usage by the guest operating system to logical processor IDs of the plurality of processors is a static mapping defined by the mapping table.

In one embodiment, the bus attachment device checks whether the copy of the device table entry is cached in a local cache operationally connected with the bus attachment device, and wherein if the copy of the device table entry is cached the retrieving of the copy of the device table entry is a retrieving from the respective cache, else the retrieving of the device table entry is a retrieving from the memory.

In one embodiment, the memory further comprises an interrupt summary vector, and the device table entry further comprises an interrupt summary vector address indicator

indicating a memory address of the interrupt summary vector, the interrupt summary vector comprising an interrupt summary indicator per bus connected module, and each interrupt summary indicator being assigned to a bus connected module indicating whether there is an interrupt signal issued by the respective bus connected module to be handled. The method further comprises using by the bus attachment device the indicated memory address of the interrupt summary vector to update the interrupt summary indicator assigned to the bus connected module from which the interrupt signal is received such that the updated interrupt summary indicator indicates that there is an interrupt signal issued by the respective bus connected module to be handled.

In one embodiment, the interrupt summary vector is implemented as a contiguous area.

In one embodiment, the interrupt summary indicators each being implemented as a single bit.

In one embodiment, the memory further comprises a directed interrupt summary vector, and the device table entry further comprises a directed interrupt summary vector address indicator indicating a memory address of the directed interrupt summary vector, the directed interrupt summary vector comprising a directed interrupt summary indicator per interrupt target ID, and each directed interrupt summary indicator being assigned to an interrupt target ID indicating whether there is an interrupt signal addressed to the respective interrupt target ID to be handled. The method further comprises using by the bus attachment device the indicated memory address of the directed interrupt summary vector to update the interrupt summary indicator assigned to the target processor ID to which the received interrupt signal is addressed such that the updated interrupt summary indicator indicates that there is an interrupt signal addressed to the respective interrupt target ID to be handled.

In one embodiment, the directed interrupt summary vector is implemented as a as contiguous area.

In one embodiment, the directed interrupt summary indicators each being implemented as a single bit.

In one embodiment, the memory further comprises one or more interrupt signal vectors, and the device table entry further comprises an interrupt signal vector address indicator indicating a memory address of an interrupt signal vector of the one or more interrupt signal vectors, each of the interrupt signal vectors comprising one or more signal indicators, and each interrupt signal indicator being assigned to a bus connected module of the one or more bus connected modules and an interrupt target ID indicating whether an interrupt signal has been received from the respective bus connected module addressed to the respective interrupt target ID. The method further comprises: using by the bus attachment device the indicated memory address of the interrupt signal vector to select the interrupt signal indicator assigned to the bus connected module which issued the received interrupt signal and to the interrupt target ID to which the received interrupt signal is addressed, and updating the selected interrupt signal indicator such that the selected interrupt signal indicator indicates that there is an interrupt signal issued by the respective bus connected module and addressed to the respective interrupt target ID to be handled.

In one embodiment, the interrupt signal vectors each comprises an interrupt signal indicator per interrupt target ID assigned to the respective interrupt target ID, each of the interrupt signal vectors being assigned to an individual bus connected module with the interrupt signal indicators of the



respective interrupt signal vector being further assigned to the respective individual bus connected module.

In one embodiment, the interrupt signal vectors each comprises an interrupt signal indicator per bus connected module assigned to the respective bus connected module, each of the interrupt signal vectors being assigned to an individual target processor ID with the interrupt signal indicators of the respective interrupt signal vector being further assigned to the respective target processor ID.

In one embodiment, the interrupt signal vectors each being implemented as a contiguous area in memory.

In one embodiment, the interrupt signal indicators each being implemented as a single bit.

In one embodiment, the device table entry further comprises a logical partition ID identifying a logical partition to which the guest operating system is assigned, and the forwarding of the interrupt signal by the bus attachment device further comprises forwarding with the interrupt signal the logical partition ID.

In one embodiment, the bus attachment module comprises a mapping table per logical partition ID.

In one embodiment, the method further comprises retrieving by the bus attachment device an interrupt subclass ID identifying an interrupt subclass to which the received interrupt signal is assigned, and the forwarding of the interrupt signal by the bus attachment device further comprises forwarding with the interrupt signal the interrupt subclass ID.

In one embodiment, the processors of the computer system are adapted to execute a plurality of guest operating systems, and the bus attachment device comprises a mapping table for each guest operating system of the plurality of guest operating systems.

In one embodiment, the method further comprises: receiving by the bus attachment device a request from the bus connected module for a direct memory access for updating status information of the bus connected module in the memory, the status update of the bus connected module triggering the interrupt signal; and upon receiving of the request, performing by the bus attachment device the direct memory access to the memory updating the status information of the bus connected module in the memory.

In one aspect, a computer system for providing an interrupt signal to a guest operating system executed using one or more processors of a plurality of processors of the computer system assigned for usage by the guest operating system is provided. The computer system further comprises one or more bus connected modules being operationally connected with the plurality of processors via a bus and a bus attachment device. Each processor of the plurality of processors is assigned with a logical processor ID used by the bus attachment device to address the respective processor; and each processor of the plurality of processors assigned for usage by the guest operating system further being assigned with an interrupt target ID used by the guest operating system and the one or more bus connected modules to address the respective processor. The computer system is configured to perform a method comprising: receiving by the bus attachment device an interrupt signal from one of the bus connected modules with an interrupt target ID, the interrupt target ID identifying one of the processors assigned for usage by the guest operating system as a target processor for handling the interrupt signal; translating by the bus attachment device the received interrupt target ID to a logical processor ID of the target processor using a mapping table comprised by the bus attachment device, the mapping table mapping the interrupt

target IDs of the processors assigned for usage by the guest operating system to logical processor IDs of the plurality of processors; and forwarding by the bus attachment device the interrupt signal to the target processor for handling using the logical processor ID of the target processor to address the target processor directly.

In one aspect, a computer program product for providing an interrupt signal to a guest operating system executed using one or more processors of a plurality of processors of a computer system assigned for usage by the guest operating system is provided. The computer system further comprises one or more bus connected modules being operationally connected with the plurality of processors via a bus and a bus attachment device. Each processor of the plurality of processors is assigned with a logical processor ID used by the bus attachment device to address the respective processor, and each processor of the plurality of processors assigned for usage by the guest operating system further being assigned with an interrupt target ID used by the guest operating system and the one or more bus connected modules to address the respective processor. The computer program product comprises a computer readable non-transitory medium readable by a processing circuit and storing instructions for execution by the processing circuit for performing a method comprising: receiving by the bus attachment device an interrupt signal from one of the bus connected modules with an interrupt target ID, the interrupt target ID identifying one of the processors assigned for usage by the guest operating system as a target processor for handling the interrupt signal; translating by the bus attachment device the received interrupt target ID to a logical processor ID of the target processor using a mapping table comprised by the bus attachment device, the mapping table mapping the interrupt target IDs of the processors assigned for usage by the guest operating system to logical processor IDs of the plurality of processors; and forwarding by the bus attachment device the interrupt signal to the target processor for handling using the logical processor ID of the target processor to address the target processor directly.

What is claimed is:

1. A computer program product for providing an interrupt signal to a guest operating system executed using one or more processors of a plurality of processors of a computer system assigned for usage by the guest operating system, the computer program product comprising:

at least one computer readable storage medium readable by at least one processing circuit and storing instructions for performing a method comprising:

receiving, by a bus attachment device of the computer system from a bus connected module of a plurality of bus connected modules operationally coupled to the plurality of processors via the bus attachment device, an interrupt signal with an interrupt target ID, the interrupt target ID identifying one processor of the plurality of processors assigned for usage by the guest operating system as a target processor for handling the interrupt signal;

translating, by the bus attachment device, the interrupt target ID to a logical processor ID of the target processor using a mapping table of the bus attachment device, the mapping table mapping interrupt target IDs of the processors assigned for usage by the guest operating system to logical processor IDs of the plurality of processors;

retrieving, by the bus attachment device, a copy of a device table entry from a device table stored in a memory operationally coupled with the bus attach-



35

ment device, the device table entry comprising a direct signaling indicator indicating whether the target processor is to be addressed directly; and

based on the direct signaling indicator indicating a direct forwarding of the interrupt signal, forwarding, by the bus attachment device, the interrupt signal to the target processor to handle, the forwarding using the logical processor ID of the target processor to address the target processor directly, else forwarding by the bus attachment device the interrupt signal for handling to the plurality of processors using broadcasting.

2. The computer program product of claim 1, wherein the interrupt signal being received is in a form of a message signaled interrupt comprising the interrupt target ID of the target processor.

3. The computer program product of claim 1, wherein the method further comprises:

checking, by the bus attachment device, whether the copy of the device table entry is cached in a local cache operationally connected with the bus attachment device;

based on the copy of the device table entry being cached, the retrieving of the copy of the device table entry is a retrieving from the cache; and

based on the copy of the device table entry not being cached, the retrieving of the device table entry is a retrieving from the memory.

4. The computer program product of claim 1, wherein the memory comprises an interrupt summary vector, the device table entry further comprises an interrupt summary vector address indicator indicating a memory address of the interrupt summary vector, the interrupt summary vector comprising an interrupt summary indicator per bus connected module, each interrupt summary indicator being assigned to a respective bus connected module indicating whether there is an interrupt signal issued by the respective bus connected module to be handled, and wherein the method further comprises:

using, by the bus attachment device, the memory address of the interrupt summary vector to update the interrupt summary indicator assigned to the bus connected module from which the interrupt signal is received such that an updated interrupt summary indicator indicates that there is the interrupt signal issued by the respective bus connected module to be handled.

5. The computer program product of claim 4, wherein the interrupt summary vector is implemented as a contiguous area.

6. The computer program product of claim 1, wherein the memory further comprises a directed interrupt summary vector, the device table entry further comprises a directed interrupt summary vector address indicator indicating a memory address of the directed interrupt summary vector, the directed interrupt summary vector comprising a directed interrupt summary indicator per interrupt target ID, each directed interrupt summary indicator being assigned to a respective interrupt target ID indicating whether there is an interrupt signal addressed to the respective interrupt target ID to be handled, and wherein the method further comprises:

using, by the bus attachment device, the memory address of the directed interrupt summary vector to update the directed interrupt summary indicator assigned to the interrupt target ID to which the received interrupt signal is addressed such that an updated directed inter-

36

rupt summary indicator indicates that there is the interrupt signal addressed to the respective interrupt target ID to be handled.

7. The computer program product of claim 1, wherein the memory further comprises one or more interrupt signal vectors, the device table entry further comprises an interrupt signal vector address indicator indicating a memory address of an interrupt signal vector of the one or more interrupt signal vectors, each of the interrupt signal vectors comprising one or more interrupt signal indicators, each interrupt signal indicator being assigned to a bus connected module of the one or more bus connected modules and to an interrupt target ID indicating whether an interrupt signal has been received from a respective bus connected module addressed to a respective interrupt target ID, and wherein the method further comprises:

using, by the bus attachment device, the memory address of the interrupt signal vector to select a selected interrupt signal indicator assigned to the bus connected module which issued the interrupt signal and to the interrupt target ID to which the interrupt signal is addressed; and

updating the selected interrupt signal indicator such that the selected interrupt signal indicator indicates that there is the interrupt signal issued by the respective bus connected module and addressed to the respective interrupt target ID to be handled.

8. The computer program product of claim 7, wherein at least the interrupt signal vector is implemented as a contiguous area.

9. The computer program product of claim 1, wherein the device table entry further comprises a logical partition ID identifying a logical partition to which the guest operating system is assigned, and wherein the forwarding of the interrupt signal by the bus attachment device further comprises forwarding the logical partition ID with the interrupt signal.

10. The computer program product of claim 1, wherein the method further comprises retrieving, by the bus attachment device, an interrupt subclass ID identifying an interrupt subclass to which the interrupt signal is assigned, and wherein the forwarding of the interrupt signal by the bus attachment device further comprises forwarding the interrupt subclass ID with the interrupt signal.

11. The computer program product of claim 1, wherein the method further comprising:

receiving, by the bus attachment device, a request from the bus connected module for a direct memory access for updating status information of the bus connected module in memory, a status update of the bus connected module triggering the interrupt signal; and

based on receiving the request, performing by the bus attachment device the direct memory access to the memory updating the status information of the bus connected module in the memory.

12. A computer system for providing an interrupt signal to a guest operating system executed using one or more processors of a plurality of processors of a computer system assigned for usage by the guest operating system, the computer system comprising:

a memory; and

a processor in communication with the memory, wherein the computer system is configured to perform a method, said method comprising:

receiving, by a bus attachment device of the computer system from a bus connected module of a plurality of bus connected modules operationally coupled to the



37

plurality of processors via the bus attachment device, an interrupt signal with an interrupt target ID, the interrupt target ID identifying one processor of the plurality of processors assigned for usage by the guest operating system as a target processor for handling the interrupt signal;

translating, by the bus attachment device, the interrupt target ID to a logical processor ID of the target processor using a mapping table of the bus attachment device, the mapping table mapping interrupt target IDs of the processors assigned for usage by the guest operating system to logical processor IDs of the plurality of processors;

retrieving, by the bus attachment device, a copy of a device table entry from a device table stored in a memory operationally coupled with the bus attachment device, the device table entry comprising a direct signaling indicator indicating whether the target processor is to be addressed directly; and

based on the direct signaling indicator indicating a direct forwarding of the interrupt signal, forwarding, by the bus attachment device, the interrupt signal to the target processor to address the target processor directly, else forwarding by the bus attachment device the interrupt signal for handling to the plurality of processors using broadcasting.

**13.** The computer system of claim **12**, wherein the method further comprises:

checking, by the bus attachment device, whether the copy of the device table entry is cached in a local cache operationally connected with the bus attachment device;

based on the copy of the device table entry being cached, the retrieving of the copy of the device table entry is a retrieving from the cache; and

based on the copy of the device table entry not being cached, the retrieving of the device table entry is a retrieving from the memory.

**14.** The computer system of claim **12**, wherein the method further comprises retrieving, by the bus attachment device, an interrupt subclass ID identifying an interrupt subclass to which the interrupt signal is assigned, and wherein the forwarding of the interrupt signal by the bus attachment device further comprises forwarding the interrupt subclass ID with the interrupt signal.

**15.** The computer system of claim **12**, wherein the method further comprising:

receiving, by the bus attachment device, a request from the bus connected module for a direct memory access for updating status information of the bus connected module in memory, a status update of the bus connected module triggering the interrupt signal; and

based on receiving the request, performing by the bus attachment device the direct memory access to the memory updating the status information of the bus connected module in the memory.

**16.** The computer system of claim **12**, wherein the memory comprises an interrupt summary vector, the device table entry further comprises an interrupt summary vector address indicator indicating a memory address of the interrupt summary vector, the interrupt summary vector comprising an interrupt summary indicator per bus connected module, each interrupt summary indicator being assigned to a respective bus connected module indicating whether there

38

is an interrupt signal issued by the respective bus connected module to be handled, and wherein the method further comprises:

using, by the bus attachment device, the memory address of the interrupt summary vector to update the interrupt summary indicator assigned to the bus connected module from which the interrupt signal is received such that an updated interrupt summary indicator indicates that there is the interrupt signal issued by the respective bus connected module to be handled.

**17.** A computer-implemented method of providing an interrupt signal to a guest operating system executed using one or more processors of a plurality of processors of a computer system assigned for usage by the guest operating system, the computer-implemented method comprising:

receiving, by a bus attachment device of the computer system from a bus connected module of a plurality of bus connected modules operationally coupled to the plurality of processors via the bus attachment device, an interrupt signal with an interrupt target ID, the interrupt target ID identifying one processor of the plurality of processors assigned for usage by the guest operating system as a target processor for handling the interrupt signal;

translating, by the bus attachment device, the interrupt target ID to a logical processor ID of the target processor using a mapping table of the bus attachment device, the mapping table mapping interrupt target IDs of the processors assigned for usage by the guest operating system to logical processor IDs of the plurality of processors;

retrieving, by the bus attachment device, a copy of a device table entry from a device table stored in a memory operationally coupled with the bus attachment device, the device table entry comprising a direct signaling indicator indicating whether the target processor is to be addressed directly; and

based on the direct signaling indicator indicating a direct forwarding of the interrupt signal, forwarding, by the bus attachment device, the interrupt signal to the target processor to handle, the forwarding using the logical processor ID of the target processor to address the target processor directly, else forwarding by the bus attachment device the interrupt signal for handling to the plurality of processors using broadcasting.

**18.** The computer-implemented method of claim **17**, further comprising:

checking, by the bus attachment device, whether the copy of the device table entry is cached in a local cache operationally connected with the bus attachment device;

based on the copy of the device table entry being cached, the retrieving of the copy of the device table entry is a retrieving from the cache; and

based on the copy of the device table entry not being cached, the retrieving of the device table entry is a retrieving from the memory.

**19.** The computer-implemented method of claim **17**, further comprising retrieving, by the bus attachment device, an interrupt subclass ID identifying an interrupt subclass to which the interrupt signal is assigned, and wherein the forwarding of the interrupt signal by the bus attachment device further comprises forwarding the interrupt subclass ID with the interrupt signal.



20. The computer-implemented method of claim 17, further comprising:

receiving, by the bus attachment device, a request from the bus connected module for a direct memory access for updating status information of the bus connected module in memory, a status update of the bus connected module triggering the interrupt signal; and

based on receiving the request, performing by the bus attachment device the direct memory access to the memory updating the status information of the bus connected module in the memory.

\* \* \* \* \*